

The McEliece Cryptosystem Resists Quantum Fourier Sampling Attacks

Hang Dinh
Indiana University South Bend
hdinh@cs.iusb.edu

Cristopher Moore
University of New Mexico
and Santa Fe Institute
moore@cs.unm.edu

Alexander Russell
University of Connecticut
acr@cse.uconn.edu

October 19, 2010

Abstract

Quantum computers can break the RSA and El Gamal public-key cryptosystems, since they can factor integers and extract discrete logarithms. If we believe that quantum computers will someday become a reality, we would like to have *post-quantum* cryptosystems which can be implemented today with classical computers, but which will remain secure even in the presence of quantum attacks.

In this article we show that the McEliece cryptosystem over *well-permuted, well-scrambled* linear codes resists precisely the attacks to which the RSA and El Gamal cryptosystems are vulnerable—namely, those based on generating and measuring coset states. This eliminates the approach of strong Fourier sampling on which almost all known exponential speedups by quantum algorithms are based. Specifically, we show that the natural case of the Hidden Subgroup Problem to which the McEliece cryptosystem reduces cannot be solved by strong Fourier sampling, or by any measurement of a coset state. We start with recent negative results on quantum algorithms for Graph Isomorphism, which are based on particular subgroups of size two, and extend them to subgroups of arbitrary structure, including the automorphism groups of linear codes. This allows us to obtain the first rigorous results on the security of the McEliece cryptosystem in the face of quantum adversaries, strengthening its candidacy for post-quantum cryptography.

1 Introduction

Considering that common public-key cryptosystems such as RSA and El Gamal are insecure against quantum attacks, the susceptibility of other well-studied public-key systems to such attacks is naturally of fundamental interest. In this article we present evidence for the strength of the McEliece cryptosystem against quantum attacks, demonstrating that the quantum Fourier sampling attacks that cripple RSA and El Gamal do not apply to the McEliece system coupled with *well-permuted*, *well-scrambled* linear codes. While our results do not rule out other quantum (or classical) attacks, they do demonstrate security against the hidden subgroup methods that have proven so powerful for computational number theory. Additionally, we partially extend results of Kempe et al. [8] concerning the subgroups of S_n reconstructible by quantum Fourier sampling.

The McEliece cryptosystem. This public-key cryptosystem was proposed by McEliece in 1978 [11], and is typically built over Goppa codes. There are two basic types of attacks known against the McEliece cryptosystem: ciphertext only attacks, and attacks on the private key. The former is unlikely to work because it relies on solving the general decoding problem, which is NP-hard. The latter can be successful on certain classes of linear codes, and is our focus. In the McEliece cryptosystem, the private key of a user Alice consists of three matrices: a $k \times n$ generator matrix M of a hidden q -ary $[n, k]$ -linear code, an invertible $k \times k$ matrix A over the finite field \mathbb{F}_q , and an $n \times n$ permutation matrix P . Both matrices A and P are selected randomly. Alice’s public key includes the matrix $M^* = AMP$, which is a generator matrix of a linear code equivalent to the secret code. An adversary may attack the private key by first computing the secret generator matrix M , and then computing¹ the secret row “scrambler” A and the secret permutation P .

There have been some successful attacks on McEliece-type public-key systems. A notable one is Sidelnikov and Shestakov’s attack [19], which efficiently computes the matrices A and MP from the public matrix AMP , in the case that the secret code is a generalized Reed-Solomon (GRS) code. Note that this attack does not reveal the secret permutation. An attack in which the secret permutation is revealed was proposed by Loidreau and Sendrier [9]. However, this attack only works with a very limited subclass of classical binary Goppa codes, namely those with a binary generator polynomial.

Although the McEliece cryptosystem is efficient and still considered (classically) secure [3], it is rarely used in practice because of the comparatively large public key (see remark 8.33 in [12]). The discovery of successful quantum attacks on RSA and El Gamal, however, have changed the landscape: as suggested by Ryan [17] and Bernstein et al. [2], the McEliece cryptosystem could become a “post-quantum” alternative to RSA.

Quantum Fourier sampling. Quantum Fourier Sampling (QFS) is a key ingredient in most efficient algebraic quantum algorithms, including Shor’s algorithms for factorization and discrete logarithm [18] and Simon’s algorithm [20]. In particular, Shor’s algorithm relies on quantum Fourier sampling over the cyclic group \mathbb{Z}_N , while Simon’s algorithm uses quantum Fourier sampling over \mathbb{Z}_2^n . In general, these algorithms solve instances of the *Hidden Subgroup Problem* (HSP) over a finite group G . Given a function f on G whose level sets are left cosets of some unknown subgroup $H < G$, i.e., such that f is constant on each left coset of H and distinct on different left cosets, they find a set of generators for the subgroup H .

The standard approach to this problem treats f as a black box and applies f to a uniform superposition over G , producing the coset state $|cH\rangle = (1/\sqrt{|H|}) \sum_{h \in H} |ch\rangle$ for a random c . We then measure $|cH\rangle$ in a

¹Recovering the secret scrambler and the secret permutation is different from the Code Equivalence problem. The former finds a transformation between two equivalent codes, while the latter decides whether two linear codes are equivalent.

Fourier basis $\{|\rho, i, j\rangle\}$ for the space $\mathbb{C}[G]$, where ρ is an irrep² of G and i, j are row and column indices of a matrix $\rho(g)$. In the *weak* form of Fourier sampling, only the representation name ρ is measured, while in the *strong* form, both the representation name and the matrix indices are measured. This produces probability distributions from which classical information can be extracted to recover the subgroup H . Moreover, since $|cH\rangle$ is block-diagonal in the Fourier basis, the optimal measurement of the coset state can always be described in terms of strong Fourier sampling.

Understanding the power of Fourier sampling in nonabelian contexts has been an ongoing project, and a sequence of negative results [5, 13, 6] have suggested that the approach is inherently limited when the underlying groups are rich enough. In particular, Moore, Russell, and Schulman [13] showed that over the symmetric group, even the strong form of Fourier sampling cannot efficiently distinguish the conjugates of most order-2 subgroups from each other or from the trivial subgroup. That is, for any $\sigma \in S_n$ with large support, and most $\pi \in S_n$, if $H = \{1, \pi^{-1}\sigma\pi\}$ then strong Fourier sampling, and therefore any measurement we can perform on the coset state, yields a distribution which is exponentially close to the distribution corresponding to $H = \{1\}$. This result implies that the GRAPH ISOMORPHISM cannot be solved by the naive reduction to strong Fourier sampling. Hallgren et al. [6] strengthened these results, demonstrating that even entangled measurements on $o(\log n!)$ coset states result in essentially information-free outcome distributions. Kempe and Shalev [7] showed that weak Fourier sampling single coset states in S_n cannot distinguish the trivial subgroup from larger subgroups H with polynomial size and non-constant minimal degree.³ They conjectured, conversely, that if a subgroup $H < S_n$ can be distinguished from the trivial subgroup by weak Fourier sampling, then the minimal degree of H must be constant. Their conjecture was later proved by Kempe, Pyber, and Shalev [8].

1.1 Our contributions

To state our results, we say that a subgroup $H < G$ is *indistinguishable by strong Fourier sampling* if the conjugate subgroups $g^{-1}Hg$ cannot be distinguished from each other or from the trivial subgroup by measuring the coset state in an arbitrary Fourier basis. A precise definition is presented in Section 3.2. Since the optimal measurement of a coset state can always be expressed as an instance of strong Fourier sampling, these results imply that no measurement of a single coset state yields any useful information about H . Based on the strategy of Moore, Russell, and Schulman [13], we first develop a general framework, formalized in Theorem 4, to determine indistinguishability of a subgroup by strong Fourier sampling. We emphasize that their results cover the case where the subgroup has order two. Our principal contribution is to show how to extend their methods to more general subgroups.

We then apply this general framework to a class of semi-direct products $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$, bounding the distinguishability for the HSP corresponding to the private-key attack on the McEliece cryptosystem, i.e., the problem of determining A and P from M^* and M . Our bound, given in Corrolary 9 of Theorem 8, depends on the minimal degree and the size of the automorphism group of the secret code, as well as on the column rank of the secret generator matrix. In particular, the rational Goppa codes have good values for these quantities, i.e., they have small automorphism groups with large minimal degree, and have generator matrices of full rank. In general, our result indicates that the McEliece cryptosystem resists all known attacks based on strong Fourier sampling if its secret q -ary $[n, k]$ -code (i) is *well-permuted*, i.e., its automorphism group has minimal degree $\Omega(n)$ and size $e^{o(n)}$, and (ii) is *well-scrambled*, i.e., it has a generator matrix of rank at least $k - o(\sqrt{n})$. Here, we assume $q^{k^2} \leq n^{0.2n}$, which implies $\log |\text{GL}_k(\mathbb{F}_q)| = O(n \log n)$, so that Alice

²Throughout the paper, we write “irrep” as short for “irreducible representation”.

³The minimal degree of a permutation group H is the minimal number of points moved by a non-identity element of H .

only needs to flip $O(n \log n)$ bits to pick a random matrix A from $\text{GL}_k(\mathbb{F}_q)$. Thus she needs only $O(n \log n)$ coin flips overall to generate her private key.

While our main application is the security of the McEliece cryptosystem, we show in addition that our general framework is applicable to other classes of groups with simpler structure, including the symmetric group and the finite general linear group $\text{GL}_2(\mathbb{F}_q)$. For the symmetric group, we extend the results of [13] to larger subgroups of S_n . Specifically, we show that any subgroup $H < S_n$ with minimal degree $m \geq \Theta(\log |H|) + \omega(\log n)$ is indistinguishable by strong Fourier sampling over S_n . In other words, if the conjugates of H can be distinguished from each other—or from the trivial subgroup—by strong Fourier sampling, then the minimum degree of H must be $O(\log |H|) + O(\log n)$. This partially extends the results of Kempe et al. [8], which apply only to weak Fourier sampling.

We go on to demonstrate another application of our general framework for the general linear group $\text{GL}_2(\mathbb{F}_q)$, giving the first negative result regarding the power of strong Fourier sampling over $\text{GL}_2(\mathbb{F}_q)$. We show that any subgroup $H < \text{GL}_2(\mathbb{F}_q)$ that does not contain non-identity scalar matrices and has order $|H| \leq q^\delta$ for some $\delta < 1/2$ is indistinguishable by strong Fourier sampling. Examples of such subgroups are those generated by a constant number of triangular unipotent matrices.

Remark Our results show that the natural reduction of McEliece to a hidden subgroup problem yields negligible information about the secret key. Thus they rule out the direct analogue of the quantum attack that breaks, for example, RSA. Our results are quite flexible in this hidden-subgroup context: they apply equally well to any HSP reduction resulting in a rich subgroup of $\text{GL}_2(\mathbb{F}_q)$, which seems to be the natural arena for the McEliece system.

Of course, our results do not rule out other quantum (or classical) attacks. Neither do they establish that a quantum algorithm for the McEliece cryptosystem would violate a natural hardness assumption, as do recent lattice cryptosystem constructions whose hardness is based on the Learning With Errors problem (e.g. Regev [15]). Nevertheless, they indicate that any such algorithm would have to use rather different ideas than those that have been proposed so far.

1.2 Summary of technical ideas

Let G be a finite group. We wish to establish general criteria for indistinguishability of subgroups $H < G$ by strong Fourier sampling. We begin with the general strategy, developed in [13], that controls the resulting probability distributions in terms of the representation-theoretic properties of G . In order to handle richer subgroups, however, we have to overcome some technical difficulties. Our principal contribution here is a “decoupling” lemma that allows us to handle the cross terms arising from pairs of nontrivial group elements.

Roughly, the approach (presented in Section 3.2) identifies two disjoint subsets, SMALL and LARGE, of irreps of G . The set LARGE consists of all irreps whose dimensions are no smaller than a certain threshold D . While D should be as large as possible, we also need to choose D small enough so that the set LARGE is large. In contrast, the representations in SMALL must have small dimension (much smaller than \sqrt{D}), and the set SMALL should be small or contain few irreps that appear in the decomposition of the tensor product representation $\rho \otimes \rho^*$ for any $\rho \in \text{LARGE}$. In addition, any irrep ρ outside SMALL must have small normalized character $|\chi_\rho(h)|/d_\rho$ for any nontrivial element $h \in H$. If there are such two sets SMALL and LARGE, and if the order of H is sufficiently small, then H is indistinguishable by strong Fourier sampling over G .

In the case $G = \text{GL}_2(\mathbb{F}_q)$, for instance, we choose SMALL as the set of all linear representations and set the threshold $D = q - 1$. The key lemma we need to prove is then that for any nonlinear irrep ρ of $\text{GL}_2(\mathbb{F}_q)$, the decomposition of $\rho \otimes \rho^*$ contains at most two inequivalent linear representations. (Lemma 11). In the

case $G = S_n$, we choose SMALL as the set Λ_c of all Young diagrams with at least $(1 - c)n$ rows or at least $(1 - c)n$ columns, and set $D = n^{dn}$, for reasonable constants $0 < c, d < 1$. For this case, we use the same techniques as in [13].

For the case $G = (\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$ corresponding to the McEliece cryptosystem, the normalized characters on the hidden subgroup K depend on the minimal degree of the automorphism group $\text{Aut}(C)$, where C is the secret code. Moreover, $|K|$ depends on $|\text{Aut}(C)|$ and the column rank of the secret generator matrix. Now we can choose SMALL as the set of all irreps constructed from tensor product representations $\tau \times \lambda$ of $\text{GL}_k(\mathbb{F}_q) \times S_n$ with $\lambda \in \Lambda_c$. Then the “small” features of Λ_c will induce the “small” features of this set SMALL. To show that any irrep outside SMALL has small normalized characters on K , we show that any Young diagram λ outside Λ_c has large dimension (Lemma 17).

2 Hidden Subgroup Attack Against McEliece Cryptosystems

2.1 An attack via the hidden shift problem

As mentioned in the Introduction, we consider the attack that involves finding the secret scrambler and permutation in a McEliece private key.

Scrambler-Permutation Problem Given two $k \times n$ generator matrices M and M^* of two equivalent linear codes over \mathbb{F}_q , the task is to find a matrix $A \in \text{GL}_k(\mathbb{F}_q)$ and an $n \times n$ permutation P matrix such that $M^* = AMP$.

The decision version of this problem, known as CODE EQUIVALENCE problem, is not easier than GRAPH ISOMORPHISM, although it is unlikely to be NP-complete [14]. The only known way to solve the Scrambler-Permutation problem using quantum Fourier sampling is to reduce it to a Hidden Shift Problem, which in turn can be reduced to a Hidden Subgroup Problem over a wreath product.

Hidden Shift Problem Let G be a finite group and Σ be some finite set. Given two functions $f_0 : G \rightarrow \Sigma$ and $f_1 : G \rightarrow \Sigma$ on G , we call an element $s \in G$ a *left shift* from f_0 to f_1 (or simply, a *shift*) if $f_0(sx) = f_1(x)$ for all $x \in G$. We are promised that there is such a shift. Find a shift.

The Scrambler-Permutation Problem is reduced to the Hidden Shift Problem over group $G = \text{GL}_k(\mathbb{F}_q) \times S_n$ by defining functions f_0 and f_1 on $\text{GL}_k(\mathbb{F}_q) \times S_n$ as follows: for all $(A, P) \in \text{GL}_k(\mathbb{F}_q) \times S_n$,

$$f_0(A, P) = A^{-1}MP, \quad f_1(A, P) = A^{-1}M^*P. \quad (1)$$

Here and from now on, we identify each $n \times n$ permutation matrix as its corresponding permutation in S_n . Apparently, $AMP = M^*$ if and only if (A^{-1}, P) is a shift from f_0 to f_1 .

2.2 Reduction from the hidden shift problem to the hidden subgroup problem

We present how to reduce the Hidden Shift Problem over group G to the HSP on the wreath product $G \wr \mathbb{Z}_2$, which can also be written as a semi-direct product $G^2 \rtimes \mathbb{Z}_2$ associated with the action of \mathbb{Z}_2 on G^2 in which the non-identity element of \mathbb{Z}_2 acts on G^2 by swapping, i.e., $1 \cdot (x, y) = (y, x)$.

Given two input functions f_0 and f_1 for a Hidden Shift Problem on G , we define the function $f : G^2 \rtimes \mathbb{Z}_2 \rightarrow \Sigma \times \Sigma$ as follows: for $(x, y) \in G^2, b \in \mathbb{Z}_2$,

$$f((x, y), b) \stackrel{\text{def}}{=} \begin{cases} (f_0(x), f_1(y)) & \text{if } b = 0 \\ (f_1(y), f_0(x)) & \text{if } b = 1 \end{cases} \quad (2)$$

We want to determine the subgroup whose cosets are distinguished by f . Recall that a function f on a group G distinguishes the right cosets of a subgroup $H < G$ if for all $x, y \in G$,

$$f(x) = f(y) \iff yx^{-1} \in H.$$

Definition. Let f be a function on a group G . We say that the function f is *injective under right multiplication* if for all $x, y \in G$,

$$f(x) = f(y) \iff f(yx^{-1}) = f(1).$$

Define the subset $G|_f$ of the group G :

$$G|_f \stackrel{\text{def}}{=} \{g \in G \mid f(g) = f(1)\}.$$

Proposition 1. Let f be a function on a group G . If f distinguishes the right cosets of a subgroup $H < G$, then f must be injective under right multiplication and $G|_f = H$. Conversely, if f is injective under right multiplication, then $G|_f$ is a subgroup and f distinguishes the right cosets of the subgroup $G|_f$.

Hence, the function f defined in (2) can distinguish the right cosets of some subgroup if and only if it is injective under right multiplication.

Lemma 2. The function f defined in (2) is injective under right multiplication if and only if f_0 is injective under right multiplication.

The proof for this lemma is straightforward on the case by case basis, so we omit it here.

Proposition 3. Assume f_0 is injective under right multiplication. Let $H_0 = G|_{f_0}$ and s be a shift. Then the function f defined in (2) distinguishes right cosets of the following subgroup of $G^2 \rtimes \mathbb{Z}_2$:

$$G^2 \rtimes \mathbb{Z}_2|_f = ((H_0, s^{-1}H_0s), 0) \cup ((H_0s, s^{-1}H_0), 1)$$

which has size $2|H_0|^2$. Recall that the set of shifts is H_0s .

To find a hidden shift from the hidden subgroup $K = G^2 \rtimes \mathbb{Z}_2|_f$, just select an element of the form $((g_1, g_2), 1)$ from K , then g_1 must belong to H_0s , which is the set of all shifts.

In the case of Scrambler-Permutation problem. Back to the Hidden Shift Problem over $G = \text{GL}_k(\mathbb{F}_q) \times S_n$ reduced from the Scrambler-Permutation problem, it is clear that the input function f_0 defined in (1) is injective under right multiplication and that

$$H_0 = \text{GL}_k(\mathbb{F}_2) \times S_n|_{f_0} = \{(A, P) \in \text{GL}_k(\mathbb{F}_2) \times S_n : A^{-1}MP = M\}.$$

3 Quantum Fourier sampling (QFS)

3.1 Preliminaries and Notation

Fix a finite group G , abelian or non-abelian, and let \widehat{G} denote the set of (complex) irreducible representations, or “irreps” for short, of G . For each irrep $\rho \in \widehat{G}$, let V_ρ denote a vector space over \mathbb{C} on which ρ acts so that ρ is a group homomorphism from G to the general linear group over V_ρ , and let d_ρ denote the dimension of V_ρ . For each ρ , we fix an orthonormal basis $B_\rho = \{\mathbf{b}_1, \dots, \mathbf{b}_{d_\rho}\}$ for V_ρ , in which we can represent each $\rho(g)$ as a $d_\rho \times d_\rho$ unitary matrix whose j^{th} column is the vector $\rho(g)\mathbf{b}_j$.

Viewing the vector space $\mathbb{C}[G]$ as the regular representation of G , we can decompose $\mathbb{C}[G]$ into irreps as the direct sum $\bigoplus_{\rho \in \widehat{G}} V_\rho^{\oplus d_\rho}$. This has a basis $\{|\rho, i, j\rangle : \rho \in \widehat{G}, 1 \leq i, j \leq d_\rho\}$, where $\{|\rho, i, j\rangle \mid 1 \leq i \leq d_\rho\}$ is a basis for the j^{th} copy of V_ρ in the decomposition of $\mathbb{C}[G]$.

Definition. The *Quantum Fourier transform* over G is the unitary operator, denoted F_G , that transforms a vector in $\mathbb{C}[G]$ from the point-mass basis $\{|g\rangle \mid g \in G\}$ into the basis given by the decomposition of $\mathbb{C}[G]$. For all $g \in G$,

$$F_G |g\rangle = \sum_{\rho, i, j} \sqrt{\frac{d_\rho}{|G|}} \rho(g)_{i,j} |\rho, i, j\rangle,$$

where $\rho(g)_{i,j}$ is the (i, j) -entry of the matrix $\rho(g)$. Alternatively, we can view $F_G |g\rangle$ as a block diagonal matrix consisting of the block $\sqrt{d_\rho/|G|} \rho(g)$ for each $\rho \in \widehat{G}$.

Notations. For each subset $X \subset G$, define $|X\rangle = (1/\sqrt{|X|}) \sum_{x \in X} |x\rangle$, which is the state of uniformly random element of X in the point-mass basis. For each $X \subset G$ and $\rho \in \widehat{G}$, define the operator

$$\Pi_X^\rho \stackrel{\text{def}}{=} \frac{1}{|X|} \sum_{x \in X} \rho(x),$$

and let $\widehat{X}(\rho)$ denote the $d_\rho \times d_\rho$ matrix block at ρ in the quantum Fourier transform of $|X\rangle$, i.e.,

$$\widehat{X}(\rho) \stackrel{\text{def}}{=} \sqrt{\frac{d_\rho}{|G||X|}} \sum_{x \in X} \rho(x) = \sqrt{\frac{d_\rho |X|}{|G|}} \Pi_X^\rho.$$

Fact. If X is a subgroup of G , then Π_X^ρ is a projection operator. That is, $(\Pi_X^\rho)^\dagger = \Pi_X^\rho$ and $(\Pi_X^\rho)^2 = \Pi_X^\rho$.

Quantum Fourier Sampling (QFS) is a standard procedure based on the Quantum Fourier Transform to solve the Hidden Subgroup Problem (HSP) (see [10] for a survey). An instance of the HSP over G consists of a black-box function $f : G \rightarrow \{0, 1\}^*$ such that $f(x) = f(y)$ if and only if x and y belong to the same left coset of H in G , for some subgroup $H \leq G$. The problem is to recover H using the oracle $O_f : |x, y\rangle \mapsto |x, y \oplus f(x)\rangle$. The general QFS procedure for this is the following:

1. Prepare a 2-register quantum state, the first in a uniform superposition of the group elements and the second with the value zero: $|\psi_1\rangle = (1/\sqrt{|G|}) \sum_{g \in G} |g\rangle |0\rangle$.
2. Query f , i.e., apply the oracle O_f , resulting in the state

$$|\psi_2\rangle = O_f |\psi_1\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle |f(g)\rangle = \frac{1}{\sqrt{|T|}} \sum_{\alpha \in T} |\alpha H\rangle |f(\alpha)\rangle$$

where T is a transversal of H in G .

3. Measure the second register of $|\psi_2\rangle$, resulting in the state $|\alpha H\rangle |f(\alpha)\rangle$ with probability $1/|T|$ for each $\alpha \in T$. The first register of the resulting state is then $|\alpha H\rangle$ for some uniformly random $\alpha \in G$.
4. Apply the quantum Fourier transform over G to the coset state $|\alpha H\rangle$ observed at step 3:

$$F_G |\alpha H\rangle = \sum_{\rho \in \widehat{G}, 1 \leq i, j \leq d_\rho} \widehat{\alpha H}(\rho)_{i,j} |\rho, i, j\rangle.$$

5. (Weak) Observe the representation name ρ . (Strong) Observe ρ and matrix indices i, j .
6. Classically process the information observed from the previous step to determine the subgroup H .

Probability distributions produced by QFS. For a particular coset αH , the probability of measuring the representation ρ in the state $F_G |\alpha H\rangle$ is

$$P_{\alpha H}(\rho) = \|\widehat{\alpha H}(\rho)\|_F^2 = \frac{d_\rho |H|}{|G|} \text{Tr}((\Pi_{\alpha H}^\rho)^\dagger \Pi_{\alpha H}^\rho) = \frac{d_\rho |H|}{|G|} \text{Tr}(\Pi_H^\rho)$$

where $\text{Tr}(A)$ denotes the trace of a matrix A , and $\|A\|_F := \sqrt{\text{Tr}(A^\dagger A)}$ is the Frobenius norm of A . The last equality is due to the fact that $\Pi_{\alpha H}^\rho = \rho(\alpha) \Pi_H^\rho$ and that Π_H^ρ is an orthogonal projector.

Since there is no point in measuring the rows [5], we are only concerned with measuring the columns. As pointed out in [13], the optimal von Neumann measurement on a coset state can always be expressed in this form for some basis B_ρ . Conditioned on observing ρ in the state $F_G |\alpha H\rangle$, the probability of measuring a given $\mathbf{b} \in B_\rho$ is $\|\widehat{\alpha H}(\rho) \mathbf{b}\|^2$. Hence the conditional probability that we observe the vector \mathbf{b} , given that we observe the representation ρ , is then

$$P_{\alpha H}(\mathbf{b} | \rho) = \frac{\|\widehat{\alpha H}(\rho) \mathbf{b}\|^2}{P_{\alpha H}(\rho)} = \frac{\|\Pi_{\alpha H}^\rho \mathbf{b}\|^2}{\text{Tr}(\Pi_{\alpha H}^\rho)} = \frac{\|\Pi_H^\rho \mathbf{b}\|^2}{\text{Tr}(\Pi_H^\rho)}$$

where in the last equality, we use the fact that as $\rho(\alpha)$ is unitary, it preserves the norm of the vector $\Pi_H^\rho \mathbf{b}$.

The coset representative α is unknown and is uniformly distributed in T . However, both distributions $P_{\alpha H}(\rho)$ and $P_{\alpha H}(\mathbf{b} | \rho)$ are independent of α and are the same as those for the state $F_G |H\rangle$. Thus, in Step 5 of the QFS procedure above, we observe $\rho \in \widehat{G}$ with probability $P_H(\rho)$, and conditioned on this event, we observe $\mathbf{b} \in B_\rho$ with probability $P_H(\mathbf{b} | \rho)$.

If the hidden subgroup is trivial, $H = \{1\}$, the conditional probability distribution on B_ρ is uniform,

$$P_{\{1\}}(\mathbf{b} | \rho) = \frac{\|\Pi_{\{1\}}^\rho \mathbf{b}\|^2}{\text{Tr}(\Pi_{\{1\}}^\rho)} = \frac{\|\mathbf{b}\|^2}{d_\rho} = \frac{1}{d_\rho}.$$

3.2 Distinguishability by QFS

We fix a finite group G and consider quantum Fourier sampling over G in the basis given by $\{B_\rho\}$. For a subgroup $H < G$ and for $g \in G$, let H^g denote the conjugate subgroup $g^{-1}Hg$. Since $\text{Tr}(\Pi_H^\rho) = \text{Tr}(\Pi_{H^g}^\rho)$, the probability distributions obtained by QFS for recovering the hidden subgroup H^g are

$$P_{H^g}(\rho) = \frac{d_\rho |H|}{|G|} \text{Tr}(\Pi_H^\rho) = P_H(\rho) \quad \text{and} \quad P_{H^g}(\mathbf{b} | \rho) = \frac{\|\Pi_{H^g}^\rho \mathbf{b}\|^2}{\text{Tr}(\Pi_H^\rho)}.$$

As $P_{H^g}(\rho)$ does not depend on g , weak Fourier sampling can not distinguish conjugate subgroups. Our goal is to point out that for certain nontrivial subgroup $H < G$, strong Fourier sampling can not efficiently distinguish the conjugates of H from each other or from the trivial one. Recall that the distribution $P_{\{1\}}(\cdot | \rho)$ obtained by performing strong Fourier sampling on the trivial hidden subgroup is the same as the uniform distribution U_{B_ρ} on the basis B_ρ . Thus, our goal can be boiled down to showing that the probability distribution $P_{H^g}(\cdot | \rho)$ is likely to be close to the uniform distribution U_{B_ρ} in total variation, for a random $g \in G$ and an irrep $\rho \in \widehat{G}$ obtained by weak Fourier sampling.

Definition. We define the *distinguishability* of a subgroup H (using strong Fourier sampling over G), denoted \mathcal{D}_H , to be the expectation of the squared L_1 -distance between $P_{H^g}(\cdot | \rho)$ and U_{B_ρ} :

$$\mathcal{D}_H \stackrel{\text{def}}{=} \mathbb{E}_{\rho, g} [\|P_{H^g}(\cdot | \rho) - U_{B_\rho}\|_1^2],$$

where ρ is drawn from \widehat{G} according to the distribution $P_H(\rho)$, and g is chosen from G uniformly at random. We say that the subgroup H is *indistinguishable* if $\mathcal{D}_H \leq \log^{-\omega(1)} |G|$.

Note that if \mathcal{D}_H is small, then the total variation distance between $P_{H^g}(\cdot | \rho)$ and U_{B_ρ} is small with high probability due to Markov's inequality: for all $\varepsilon > 0$,

$$\Pr_g [\|P_{H^g}(\cdot | \rho) - U_{B_\rho}\|_{t.v.} \geq \varepsilon/2] = \Pr_g [\|P_{H^g}(\cdot | \rho) - U_{B_\rho}\|_1^2 \geq \varepsilon^2] \leq \mathcal{D}_H / \varepsilon^2.$$

In particular, if the subgroup H is indistinguishable by strong Fourier sampling, then for all constant $c > 0$,

$$\|P_{H^g}(\cdot | \rho) - U_{B_\rho}\|_{t.v.} < \log^{-c} |G|$$

with probability at least $1 - \log^{-c} |G|$ in both g and ρ . Indeed, our notion of indistinguishability is inspired by that of Kempe and Shalev [7]. Focusing on weak Fourier sampling, they say that H is indistinguishable if $\|P_H(\cdot) - P_{\{1\}}(\cdot)\|_{t.v.} < \log^{-\omega(1)} |G|$.

Our main theorem below will serve as a general guideline for bounding the distinguishability of H . For this bound, we define, for each $\sigma \in \widehat{G}$, the *maximal normalized character of σ on H* as

$$\overline{\chi}_\sigma(H) \stackrel{\text{def}}{=} \max_{h \in H \setminus \{1\}} \frac{|\chi_\sigma(h)|}{d_\sigma}.$$

For each subset $S \subset \widehat{G}$, let

$$\overline{\chi}_S(H) = \max_{\sigma \in \widehat{G} \setminus S} \overline{\chi}_\sigma(H) \quad \text{and} \quad d_S = \max_{\sigma \in S} d_\sigma.$$

In addition, for each reducible representation ρ of G , we let $I(\rho)$ denote the set of irreps of G that appear in the decomposition of ρ into irreps.

Theorem 4. (MAIN THEOREM) Suppose S is a subset of \widehat{G} . Let $D > d_S^2$ and $L = L_D \subset \widehat{G}$ be the set of all irreps of dimension at least D . Let

$$\Delta = \Delta_{S,L} = \max_{\rho \in L} |S \cap I(\rho \otimes \rho^*)|. \tag{3}$$

Then the distinguishability of H is bounded by

$$\mathcal{D}_H \leq 4|H|^2 \left(\overline{\chi}_S(H) + \Delta \frac{d_S^2}{D} + \frac{|\overline{L}|D^2}{|G|} \right).$$

Intuitively, the set S consists of irreps of small dimension, and L consists of irreps of large dimension. Moreover, we wish to have that the size of S is small while the size of L is large, so that most irreps are likely in L . In the cases where there are relatively few irreps, i.e. $|S| \ll D$ and $|\widehat{G}| \ll |G|$, we can simply upper bound Δ by $|S|$ and upper bound $|\widehat{L}|$ by $|\widehat{G}|$.

We discuss the proof of this theorem in Section 5.

4 Applications

In this section, we point out some applications of Theorem 4 to analyze strong Fourier sampling over certain non-abelian groups.

4.1 Strong Fourier sampling over S_n

In this part, we consider the case where G is the symmetric group S_n . Recall that each irrep of S_n is one-to-one corresponding to an integer partition $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_t)$ of n , which is associated with the Young diagram of t rows in which the i^{th} row contains λ_i columns. The conjugate representation of λ is the irrep corresponding to the partition $\lambda' = (\lambda'_1, \lambda'_2, \dots, \lambda'_{t'})$, which is obtained by flipping the Young diagram λ about the diagonal. In particular, $\lambda'_1 = t$ and $t' = \lambda_1$.

As in [13], we use Roichman's upper bound [16] on normalized characters.

Theorem 5 (Roichman's Theorem [16]). *There exist constant $b > 0$ and $0 < q < 1$ so that for $n > 4$, for every $\pi \in S_n$, and for every irrep λ of S_n ,*

$$\left| \frac{\chi_\lambda(\pi)}{d_\lambda} \right| \leq \left(\max \left(q, \frac{\lambda_1}{n}, \frac{\lambda'_1}{n} \right) \right)^{b \cdot \text{supp}(\pi)}$$

where $\text{supp}(\pi) = \#\{k \in [n] \mid \pi(k) \neq k\}$ is the support of π .

This bound works well for unbalanced Young diagrams. In particular, for a constant $0 < c < 1/4$, let Λ_c denote the collection of partitions λ of n with the property that either $\frac{\lambda_1}{n} \geq 1 - c$ or $\frac{\lambda'_1}{n} \geq 1 - c$, i.e., the Young diagram λ contains at least $(1 - c)n$ rows or contains at least $(1 - c)n$ columns. Then, Roichman's upper bound implies that for every $\pi \in S_n$ and $\lambda \notin \Lambda_c$, and a universal constant $\alpha > 0$,

$$\left| \frac{\chi_\lambda(\pi)}{d_\lambda} \right| \leq e^{-\alpha \cdot \text{supp}(\pi)}. \quad (4)$$

On the other hand, both $|\Lambda_c|$ and the maximal dimension of representations in Λ_c are small, as shown in the following Lemma of [13].

Lemma 6 (Lemma 6.2 in [13]). *Let $p(n)$ denote the number of integer partitions of n . Then $|\Lambda_c| \leq 2cn \cdot p(cn)$, and $d_\mu < n^{cn}$ for any $\mu \in \Lambda_c$.*

To give a more concrete bound for the size of Λ_c , we record the asymptotic formula for the partition function $p(n)$ [4, pg. 45]: $p(n) \approx e^{\pi\sqrt{2n/3}} / (4\sqrt{3}n) = e^{O(\sqrt{n})} n^{-1}$ as $n \rightarrow \infty$.

Now we are ready to prove the main result of this section, which is another application of Theorem 4.

Theorem 7. *Let H be a nontrivial subgroup of S_n with minimal degree m , i.e., $m = \min_{\pi \in H \setminus \{1\}} \text{supp}(\pi)$. Then for sufficiently large n , $\mathcal{D}_H \leq O(|H|^2 e^{-\alpha m})$.*

Proof. Let $2c < d < 1/2$ be constants. We will apply Theorem 4 by setting $S = \Lambda_c$ and $D = n^{dn}$. The condition $2c < d$ guarantees that $D > d_S^2$, since $d_S \leq n^{cn}$ by Lemma 6.

First, we need to bound the maximal normalized character $\overline{\chi}_S(H)$. By (4), we have $\overline{\chi}_\mu(H) \leq e^{-\alpha m}$ for all $\mu \in \widehat{S}_n \setminus S$. Hence, $\overline{\chi}_S(H) \leq e^{-\alpha m}$.

To bound the second term in the upper bound of Theorem 4, as $\Delta \leq |S|$, it suffices to bound:

$$\begin{aligned} |S| \cdot \frac{d_S^2}{D} &\leq 2cn \cdot p(cn) \cdot \frac{n^{2cn}}{n^{dn}} && \text{(by Lemma 6)} \\ &\leq e^{O(\sqrt{n})} \cdot n^{(2c-d)n} && \text{(since } cn \cdot p(cn) = e^{O(\sqrt{n})} \text{)} \\ &\leq n^{-\gamma n}/2 && \text{for sufficiently large } n, \text{ so long as } \gamma < d - 2c. \end{aligned}$$

Now bounding the last term in the upper bound of Theorem 4:

$$\begin{aligned} \frac{|\overline{L}_D| D^2}{|S_n|} &\leq \frac{p(n) n^{2dn}}{n!} && \text{(since } |\overline{L}_D| \leq |\widehat{S}_n| = p(n)) \\ &\leq \frac{e^{O(\sqrt{n})} n^{2dn}}{n^n e^{-n}} && (n! > n^n e^{-n} \text{ by Stirling's approximation)} \\ &\leq e^{O(n)} n^{(2d-1)n} \\ &\leq n^{-\gamma n}/2 && \text{for sufficiently large } n, \text{ so long as } \gamma < 1 - 2d. \end{aligned}$$

By Theorem 4, $\mathcal{D}_H \leq 4|H|^2(e^{-\alpha m} + n^{-\gamma n})$.

□

Theorem 7 generalizes Moore, Russell, and Schulman's result [13] on strong Fourier sampling over S_n , which only applied in the case $|H| = 2$. To relate our result to the results of Kempe et al. [8] observe that, since $\log |S_n| = \Theta(n \log n)$, the subgroup H is indistinguishable by strong Fourier sampling if $|H|^2 e^{-\alpha m} \leq (n \log n)^{-\omega(1)}$, or equivalently, if $m \geq (2/\alpha) \log |H| + \omega(\log n)$.

4.2 Strong Fourier sampling and the McEliece cryptosystem

Our main application of Theorem 4 is to show the limitations of strong Fourier sampling in attacking the McEliece cryptosystem. Throughout this section, we fix system parameters n, k, q of the McEliece cryptosystem, and fix a $k \times n$ generator matrix M in a private-key of the system. Recall that the known possible quantum attack against this McEliece cryptosystem involves solving the HSP over the wreath product group $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$ with the hidden subgroup being

$$K = ((H_0, s^{-1}H_0s), 0) \cup ((H_0s, s^{-1}H_0), 1) \quad (5)$$

for some hidden element $s \in \text{GL}_k(\mathbb{F}_q) \times S_n$. Here, H_0 is a subgroup of $\text{GL}_k(\mathbb{F}_q) \times S_n$ given by

$$H_0 = \{(A, P) \in \text{GL}_k(\mathbb{F}_2) \times S_n : A^{-1}MP = M\}. \quad (6)$$

Let $\text{Aut}(M)$ denote the automorphism group of the linear code generated by M . Observe that every element $(A, P) \in H_0$ must have $P \in \text{Aut}(M)$. This allows us to control the maximal normalized characters on K through the minimal degree of $\text{Aut}(M)$. Then applying Theorem 4, we show that

Theorem 8. Assume $q^{k^2} \leq n^{an}$ for some constant $0 < a < 1/4$. Let m be the minimal degree of the automorphism group $\text{Aut}(M)$. Then for sufficiently large n , the subgroup K defined in (5) has $\mathcal{D}_K \leq O(|K|^2 e^{-\delta m})$, where $\delta > 0$ is a constant.

The proof of Theorem 8 follows the technical ideas discussed in the Introduction. The details appear in Section 6. As $q^{k^2} \leq n^{an}$, we have $\log |(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2| = O(\log n! + \log q^{k^2}) = O(n \log n)$. Hence, the subgroup K is indistinguishable if $|K|^2 e^{-\delta m} \leq (n \log n)^{-\omega(1)}$.

In general, the size of the subgroup K depends on the size of the automorphism group $\text{Aut}(M)$ and the column rank of the matrix M . To see this, we have $|K| = 2|H_0|^2$, and $|H_0| = |\text{Aut}(M)| \times |\text{Fix}(M)|$, where $\text{Fix}(M)$ denotes the set of matrices in $\text{GL}_k(\mathbb{F}_q)$ fixing M , i.e.,

$$\text{Fix}(M) = \{S \in \text{GL}_k(\mathbb{F}_q) \mid SM = M\}.$$

To bound the size of $\text{Fix}(M)$, we record an easy fact which can be obtained by the orbit-stabilizer formula.

Fact. Let r be the column rank of M . Then $|\text{Fix}(M)| = (q^k - q^r)(q^k - q^{r+1}) \dots (q^k - q^{k-1}) \leq q^{k(k-r)}$.

Proof. WLOG, assume the first r columns of M are linearly independent, and each remaining column is a linear combination of the first r columns. Consider the action of $\text{GL}_k(\mathbb{F}_q)$ on the set of $k \times n$ matrices over \mathbb{F}_q . Under this action, the orbit of the matrix M , denoted $\text{Orb}(M)$, consists of all $k \times m$ matrices over \mathbb{F}_q such that the first r columns are linearly independent, and each j^{th} column, for $j > r$, consists of the same linear combination of the first r columns as that of the j^{th} column of the matrix M . Hence, the size of $\text{Orb}(M)$ equals the number of $k \times r$ matrices over \mathbb{F}_q of column rank r . Thus, $\text{Orb}(M) = (q^k - 1)(q^k - q) \dots (q^k - q^{r-1})$. On the other hand, $\text{Fix}(M)$ is the stabilizer of M . By the orbit-stabilizer formula, we have

$$|\text{Fix}(M)| = \frac{|\text{GL}_k(\mathbb{F}_q)|}{|\text{Orb}(M)|} = \frac{(q^k - 1)(q^k - q) \dots (q^k - q^{k-1})}{(q^k - 1)(q^k - q) \dots (q^k - q^{r-1})} = (q^k - q^r)(q^k - q^{r+1}) \dots (q^k - q^{k-1}).$$

□

Corollary 9. Assume $q^{k^2} \leq n^{0.2n}$ and the automorphism group $\text{Aut}(M)$ has minimal degree $\Omega(n)$. Let r be the column rank of M . Then the subgroup K defined in (5) has $\mathcal{D}_K \leq |\text{Aut}(M)|^4 q^{4k(k-r)} e^{-\Omega(n)}$. In particular, the subgroup K is indistinguishable if, further, $|\text{Aut}(M)| \leq e^{o(n)}$ and $r \geq k - o(\sqrt{n})$.

In the a case that the matrix M generates a rational Goppa code, then M has full rank and the automorphism group $\text{Aut}(M)$ is isomorphic to a subgroup of the projective linear group $\text{PGL}_2(\mathbb{F}_q)$, provided $2 \leq k \leq n - 2$, by Stichtenoth's Theorem [21] (see Appendix B for more detailed background on rational Goppa codes). This important property results in very good values we could desire for the automorphism group $\text{Aut}(M)$: we have $|\text{Aut}(M)| \leq |\text{PGL}_2(\mathbb{F}_q)| \leq q^3$, and moreover,

Lemma 10. If M generates a rational Goppa code, the minimal degree of $\text{Aut}(M)$ is at least $n - 3$.

Proof. (sketch) Since $\text{Aut}(M)$ is isomorphic to a subgroup of $\text{PGL}_k(\mathbb{F}_q)$, the proof is based on the observation that any transformation in $\text{PGL}_k(\mathbb{F}_q)$ that fixes at least three distinct projective lines must be the identity. □

By Corollary 9, the McEliece cryptosystem coupled with rational Goppa codes resists known quantum attacks based on strong Fourier sampling. Unfortunately, this cryptosystem is insecure due to the classical attack by Sidelnikov and Shestakov [19] that recovers the secret scrambler A and the product MP , but does not reveal the secret permutation P . Of course, our next goal will be to find other classes of linear codes with which the McEliece cryptosystem would be secure against both classical and quantum attacks.

4.3 Strong Fourier sampling over $\text{GL}_2(\mathbb{F}_q)$

In this simple application, we consider the finite general linear group $G = \text{GL}_2(\mathbb{F}_q)$, whose structure as well as irreps are well established [4, §5.2]. From the character table of $\text{GL}_2(\mathbb{F}_q)$, which can be found in Appendix C, we draw the following easy facts:

Fact. Let σ be an irrep of $\text{GL}_2(\mathbb{F}_q)$. Then (i) For all $g \in \text{GL}_2(\mathbb{F}_q)$, $|\chi_\sigma(g)| = d_\sigma$ if g is a scalar matrix, and $|\chi_\sigma(g)| \leq 2$ otherwise. (ii) If $d_\sigma > 1$, then $q - 1 \leq d_\sigma \leq q + 1$.

Let H be a subgroup of $\text{GL}_2(\mathbb{F}_q)$. If H contains a non-identity scalar matrix, we have $\overline{\chi}_\sigma(H) = 1$ for all σ , making it impossible to find a set of irreps whose maximal normalized characters on H are small enough to apply our general theorem (Theorem 4). For this reason, we shall assume that H does not contain scalar matrices except for the identity. An example of such a subgroup H is any group lying inside the subgroup of triangular unipotent matrices $\{T(b) \mid b \in \mathbb{F}_q\}$, where $T(b) := \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix}$.

From the easy facts above for $\text{GL}(2, q)$, it is natural to choose the set S in Theorem 4 to be the set of linear (i.e., 1-dimensional) representations, and choose the dimensional threshold D to be $q - 1$. However, since $\text{GL}(2, q)$ has $q - 1$ linear representations, i.e., $|S| = D$, we can't upper bound Δ by $|S|$. We prove the following lemma to provide a strong upper bound on Δ , which is, in this case, the maximal number of linear representations appearing in the decomposition of $\rho \otimes \rho^*$, for any nonlinear irrep ρ .

Lemma 11. *Let ρ be an irrep of $\text{GL}(2, q)$. Then at most two linear representations appear in the decomposition of $\rho \otimes \rho^*$.*

The proof for this lemma can be found in Appendix C. Then applying Theorem 4 with S being the set of linear representations, and L being the set of non-linear irreps of $\text{GL}_2(\mathbb{F}_q)$, we have:

Corollary 12. *Let H be a subgroup of $\text{GL}_2(\mathbb{F}_q)$ that does not contain any scalar matrix other than the identity. Then $\mathcal{D}_H \leq 28|H|^2/q$.*

Proof of Corollary 12. Let S be the set of linear representations of $\text{GL}_2(\mathbb{F}_q)$ and let $D = q - 1$. Then in this case, L_D is the set of all non-linear irreps of $\text{GL}_2(\mathbb{F}_q)$.

Since $\overline{\chi}_\sigma(H) \leq 2/(q - 1)$ for all nonlinear irrep σ , we have

$$\overline{\chi}_{\overline{S}}(H) \leq 2/(q - 1) \leq 0.5/|H|.$$

To bound the second term in the bound of 4, we have $\Delta \leq 2$ by Lemma 11 and $d_S = 1$, thus

$$\Delta \frac{d_S^2}{D} \leq 2/(q - 1) \leq 3/q.$$

As $|G| = (q - 1)^2 q(q + 1)$ and $|\overline{L}_D| = |S| = q - 1$, we have

$$\frac{|\overline{L}_D| D^2}{|G|} = \frac{(q - 1)^3}{(q - 1)^2 q(q + 1)} = \frac{q - 1}{q(q + 1)} < 1/q.$$

By Theorem 4, $\mathcal{D}_H \leq 4|H|^2 (7/q)$. □

In particular, H is indistinguishable by strong Fourier sampling over $\text{GL}_2(\mathbb{F}_q)$ if $|H| \leq q^\delta$ for some $\delta < 1/2$, because in that case we have $\mathcal{D}_H \leq 28q^{2\delta-1} \leq \log^{-c} |\text{GL}_2(\mathbb{F}_q)|$ for all constant $c > 0$.

Examples of indistinguishable subgroups. As a specific example, consider a cyclic subgroup H_b generated by a triangular unipotent matrix $T(b)$ for any $b \neq 0$. Since $T(b)^k = T(kb)$ for any integer $k \geq 0$, the order of H_b is the least positive integer k such that $kb = 0$. In particular, the order of H_b equals the characteristic of the finite field \mathbb{F}_q . Suppose $q = p^n$ for some prime number p and $n > 2$. Then \mathbb{F}_q has characteristic p , and hence, $|H_b| = p$. By Corollary 12, we have $\mathcal{D}_{H_b} \leq O(p^{2-n})$.

Similarly, consider a subgroup $H_{a,b}$ generated by two distinct non-identity elements $T(a)$ and $T(b)$. Since elements of $H_{a,b}$ are of the form $T(ka + \ell b)$ for $k, \ell \in \{0, 1, \dots, p-1\}$, we have $|H_{a,b}| \leq p^2$. Thus, the distinguishability of $H_{a,b}$ using strong Fourier sampling over $\text{GL}_2(\mathbb{F}_{p^n})$ is $O(p^{4-n})$. Clearly, both H_b and $H_{a,b}$ are indistinguishable, for n sufficiently large. More generally, any subgroup generated by a constant number of triangular unipotent matrices is indistinguishable.

5 Bounding distinguishability

We now present the proof for the main theorem (Theorem 4) in details. Fixing a nontrivial subgroup $H < G$, we want to upper bound \mathcal{D}_H . Let us start with bounding the expectation over the random group element $g \in G$, for a fixed irrep $\rho \in \widehat{G}$:

$$E_H(\rho) \stackrel{\text{def}}{=} \mathbb{E}_g [\|P_{H^g}(\cdot | \rho) - U_{B_\rho}\|_1^2].$$

Obviously we always have $E_H(\rho) \leq 4$. More interestingly, we have

$$\begin{aligned} E_H(\rho) &= \mathbb{E}_g \left[\left(\sum_{\mathbf{b} \in B_\rho} \left| P_{H^g}(\mathbf{b} | \rho) - \frac{1}{d_\rho} \right| \right)^2 \right] \\ &\leq \mathbb{E}_g \left[d_\rho \sum_{\mathbf{b} \in B_\rho} \left(P_{H^g}(\mathbf{b} | \rho) - \frac{1}{d_\rho} \right)^2 \right] \quad (\text{by Cauchy-Schwarz}) \\ &= d_\rho \sum_{\mathbf{b} \in B_\rho} \text{Var}_g [P_{H^g}(\mathbf{b} | \rho)] \quad (\text{since } \mathbb{E}_g [P_{H^g}(\mathbf{b} | \rho)] = \frac{1}{d_\rho}) \\ &= \frac{d_\rho}{\text{Tr}(\Pi_H^\rho)^2} \sum_{\mathbf{b} \in B_\rho} \text{Var}_g [\|\Pi_{H^g}^\rho \mathbf{b}\|^2]. \end{aligned} \tag{7}$$

The equation $\mathbb{E}_g [P_{H^g}(\mathbf{b} | \rho)] = 1/d_\rho$ (Proposition 18 in Appendix A) can be shown using *Schur's lemma*.

From (7), we are motivated to bound the variance of $\|\Pi_{H^g}^\rho \mathbf{b}\|^2$ when g is chosen uniformly at random. We provide an upper bound that depends on the projection of the vector $\mathbf{b} \otimes \mathbf{b}^*$ onto irreducible subspaces of $\rho \otimes \rho^*$, and on maximal normalized characters of σ on H for all irreps σ appearing in the decomposition of $\rho \otimes \rho^*$. Recall that the representation $\rho \otimes \rho^*$ is typically reducible and can be written as an orthogonal direct sum of irreps $\rho \otimes \rho^* = \bigoplus_{\sigma \in \widehat{G}} a_\sigma \sigma$, where $a_\sigma \geq 0$ is the multiplicity of σ . Then $I(\rho \otimes \rho^*)$ consists of σ with $a_\sigma > 0$, and we let $\Pi_\sigma^{\rho \otimes \rho^*}$ denote the projection operator whose image is $a_\sigma \sigma$, that is, the subspace spanned by all copies of σ . Our upper bound given in Lemma 13 below generalizes the bound given in Lemma 4.3 of [13], which only applies to subgroups H of order 2.

Lemma 13. *Let ρ be an irrep of G . Then for any vector $\mathbf{b} \in V_\rho$,*

$$\text{Var}_g [\|\Pi_{H^g}^\rho \mathbf{b}\|^2] \leq \sum_{\sigma \in I(\rho \otimes \rho^*)} \overline{\chi}_\sigma(H) \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2.$$

Proof of Lemma 13. Fix a vector $\mathbf{b} \in V_\rho$. To simplify notations, we shall write Π_g as shorthand for $\Pi_{H^g}^\rho$, and write $g\mathbf{b}$ for $\rho(g)\mathbf{b}$. For any $g \in G$, we have

$$\begin{aligned}\|\Pi_g \mathbf{b}\|^2 &= \langle \Pi_g \mathbf{b}, \Pi_g \mathbf{b} \rangle = \langle \mathbf{b}, \Pi_g \mathbf{b} \rangle \\ &= \frac{1}{|H|} \left(\langle \mathbf{b}, \mathbf{b} \rangle + \sum_{h \in H \setminus \{1\}} \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle \right).\end{aligned}$$

Let $S_g = \sum_{h \in H \setminus \{1\}} \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle$. Then

$$\|\Pi_g \mathbf{b}\|^2 = \frac{1}{|H|} (\|\mathbf{b}\|^2 + S_g) \quad \text{and} \quad S_g = |H| \|\Pi_g \mathbf{b}\|^2 - \|\mathbf{b}\|^2.$$

It follows that S_g is real, and that

$$\|\Pi_g \mathbf{b}\|^4 = \frac{1}{|H|^2} (\|\mathbf{b}\|^4 + 2\|\mathbf{b}\|^2 S_g + S_g^2).$$

We have

$$\mathbb{E}_g [\|\Pi_g \mathbf{b}\|^4] = \frac{1}{|H|^2} (\|\mathbf{b}\|^4 + 2\|\mathbf{b}\|^2 \mathbb{E}_g [S_g] + \mathbb{E}_g [S_g^2]) \quad (8)$$

$$\begin{aligned}\mathbb{E}_g [\|\Pi_g \mathbf{b}\|^2]^2 &= \frac{1}{|H|^2} (\|\mathbf{b}\|^2 + \mathbb{E}_g [S_g])^2 \\ &= \frac{1}{|H|^2} (\|\mathbf{b}\|^4 + 2\|\mathbf{b}\|^2 \mathbb{E}_g [S_g] + \mathbb{E}_g [S_g]^2)\end{aligned} \quad (9)$$

Subtracting (8) by (9) yields

$$\begin{aligned}\text{Var}_g [\|\Pi_g \mathbf{b}\|^2] &= \mathbb{E}_g [\|\Pi_g \mathbf{b}\|^4] - \mathbb{E} [\|\Pi_g \mathbf{b}\|^2]^2 \\ &= \frac{1}{|H|^2} (\mathbb{E}_g [S_g^2] - \mathbb{E}_g [S_g]^2) \leq \frac{1}{|H|^2} \mathbb{E}_g [S_g^2].\end{aligned}$$

To bound the variance, we upper bound S_g^2 for all $g \in G$. Since S_g is real, applying Cauchy-Schwarz inequality, we have

$$S_g^2 = \left| \sum_{h \in H \setminus \{1\}} \langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle \right|^2 \leq (|H| - 1) \left(\sum_{h \in H \setminus \{1\}} |\langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle|^2 \right).$$

Proving similarly to Lemma 4.2 in [13], one can express the second moment of the inner product $\langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle$ in terms of the projection of $\mathbf{b} \otimes \mathbf{b}^*$ into the irreducible constituents of the tensor product representation $\rho \otimes \rho^*$. Specifically, for any $h \in G$, we have

$$\mathbb{E}_g [|\langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle|^2] = \sum_{\sigma \in I(\rho \otimes \rho^*)} \frac{\chi_\sigma(h)}{d_\sigma} \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2.$$

It follows that

$$\text{Var}_g [\|\Pi_{H^g}^\rho \mathbf{b}\|^2] \leq \frac{1}{|H|} \left(\sum_{h \in H \setminus \{1\}} \mathbb{E}_g [|\langle \mathbf{b}, g^{-1} h g \mathbf{b} \rangle|^2] \right) \leq \sum_{\sigma \in I(\rho \otimes \rho^*)} \overline{\chi}_\sigma(H) \left\| \Pi_\sigma^{\rho \otimes \rho^*} (\mathbf{b} \otimes \mathbf{b}^*) \right\|^2.$$

□

Back to our goal of bounding $E_H(\rho)$ using the bound in Lemma 13, the strategy will be to separate irreps appearing in the decomposition of $\rho \otimes \rho^*$ into two groups, those with small dimension and those with large dimension, and treat them differently. If d_σ is large, we shall rely on bounding $\overline{\chi}_\sigma(H)$. If d_σ is small, we shall control the projection given by $\Pi_\sigma^{\rho \otimes \rho^*}$ using the following lemma which was proved implicitly in [13] (its proof is also given in Appendix):

Lemma 14. *For any irrep σ , we have $\sum_{\mathbf{b} \in B_\rho} \left\| \Pi_\sigma^{\rho \otimes \rho^*}(\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 \leq d_\sigma^2$.*

The method discussed above for bounding $E_H(\rho)$ is culminated into Lemma 15 below.

Lemma 15. *Let $\rho \in \widehat{G}$ be arbitrary and $S \subset \widehat{G}$ be any subset of irreps that does not contain ρ . Then*

$$E_H(\rho) \leq 4|H|^2 \left(\overline{\chi}_{\overline{S}}(H) + |S \cap I(\rho \otimes \rho^*)| \frac{d_S^2}{d_\rho} \right).$$

Proof of Lemma 15. Combining Inequality (7) and Lemmas 13 give

$$E_H(\rho) \leq \frac{d_\rho}{\text{Tr}(\Pi_H^\rho)^2} \sum_{\sigma \in I(\rho \otimes \rho^*)} \overline{\chi}_\sigma(H) \sum_{\mathbf{b} \in B_\rho} \left\| \Pi_\sigma^{\rho \otimes \rho^*}(\mathbf{b} \otimes \mathbf{b}^*) \right\|^2.$$

Now we split additive items in the above upper bound into two groups separated by the set S . For the first group (large dimension),

$$\begin{aligned} \sum_{\sigma \in \overline{S} \cap \widehat{G}^{\rho \otimes \rho^*}} \overline{\chi}_\sigma(H) \sum_{\mathbf{b} \in B_\rho} \left\| \Pi_\sigma^{\rho \otimes \rho^*}(\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 &\leq \overline{\chi}_{\overline{S}}(H) \sum_{\mathbf{b} \in B_\rho} \underbrace{\sum_{\sigma \in I(\rho \otimes \rho^*)} \left\| \Pi_\sigma^{\rho \otimes \rho^*}(\mathbf{b} \otimes \mathbf{b}^*) \right\|^2}_{\leq 1} \\ &\leq \overline{\chi}_{\overline{S}}(H) d_\rho. \end{aligned}$$

For the second group (small dimension),

$$\begin{aligned} \sum_{\sigma \in S \cap I(\rho \otimes \rho^*)} \overline{\chi}_\sigma(H) \sum_{\mathbf{b} \in B_\rho} \left\| \Pi_\sigma^{\rho \otimes \rho^*}(\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 &\leq \sum_{\sigma \in S \cap I(\rho \otimes \rho^*)} \overline{\chi}_\sigma(H) d_\sigma^2 \quad (\text{by Lemma 14}) \\ &\leq \sum_{\sigma \in S \cap I(\rho \otimes \rho^*)} d_\sigma^2 \quad (\text{since } \overline{\chi}_\sigma(H) \leq 1) \\ &\leq |S \cap I(\rho \otimes \rho^*)| d_S^2. \end{aligned}$$

Summing the last bounds for the two groups yields

$$E_H(\rho) \leq \left(\frac{d_\rho}{\text{Tr}(\Pi_H^\rho)} \right)^2 \left(\overline{\chi}_{\overline{S}}(H) + |S \cap I(\rho \otimes \rho^*)| \frac{d_S^2}{d_\rho} \right).$$

On the other hand, since $E_H(\rho) \leq 4$, we can assume $H^2 \overline{\chi}_{\overline{S}}(H) \leq 1$, and thus $\overline{\chi}_{\overline{S}}(H) \leq \frac{1}{|H|^2} \leq \frac{1}{2|H|}$. Hence, we have

$$\frac{\text{Tr}(\Pi_H^\rho)}{d_\rho} = \frac{1}{|H|} \left(1 + \sum_{h \in H \setminus \{1\}} \frac{\chi_\rho(h)}{d_\rho} \right) \geq \frac{1}{|H|} - \overline{\chi}_\rho(H) \geq \frac{1}{2|H|},$$

where the last inequality is due to $\overline{\chi}_\rho(H) \leq \overline{\chi}_{\overline{S}}(H) \leq \frac{1}{2|H|}$. This completes the proof. \square

To apply this lemma, we should choose the subset S such that $d_S^2 \ll d_\rho$, that is, S should consist of small dimensional irreps. Then applying Lemma 15 for all irreps ρ of large dimension, we can prove our general main theorem straightforwardly.

Proof of Theorem 4: For any $\rho \in L$, since $d_\rho \geq D > d_S^2$, we must have $\rho \notin S$. By Lemma 15,

$$E_H(\rho) \leq 4|H|^2 \left(\overline{\chi}_S(H) + \Delta \frac{d_S^2}{D} \right) \quad \text{for all } \rho \in L.$$

Combining this with the fact that $E_H(\rho) \leq 4$ for all $\rho \notin L$, we obtain

$$\mathcal{D}_H = \mathbb{E}_\rho[E_H(\rho)] \leq 4|H|^2 \left(\overline{\chi}_S(H) + \Delta \frac{d_S^2}{D} \right) + 4\Pr_\rho[\rho \notin L].$$

To complete the proof, it remains to bound $\Pr_\rho[\rho \notin L]$. Since $\text{Tr}(\Pi_H^\rho) \leq d_\rho$, we have

$$P(\rho) = \frac{d_\rho |H|}{|G|} \text{Tr}(\Pi_H^\rho) \leq \frac{d_\rho^2 |H|}{|G|}.$$

Since $d_\rho < D$ for all $\rho \in \widehat{G} \setminus L$, it follows that

$$\Pr_\rho[\rho \notin L] = \sum_{\rho \notin L} P(\rho) \leq \frac{|\overline{L}| D^2 |H|}{|G|} \leq \frac{|\overline{L}| D^2 |H|^2}{|G|}.$$

□

6 Strong Fourier sampling over $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$

This section devotes to the proof of Theorem 8 which establishes the limitation of strong Fourier sampling in breaking the McEliece cryptosystem. The goal is to bound the distinguishability of the subgroup K defined in (5) of the wreath product $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$.

6.1 Normalized characters for $G \wr \mathbb{Z}_2$

Firstly, we consider quantum Fourier sampling over the wreath product $G \wr \mathbb{Z}_2$, for a general group G , with a hidden subgroup of the form

$$K = ((H_0, s^{-1}H_0s), 0) \cup ((H_0s, s^{-1}H_0), 1) < G \wr \mathbb{Z}_2$$

for some subgroup $H_0 < G$ and some element $s \in G$. Again, the first thing we need to understand is the maximal normalized characters on K . Recall that all irreducible characters of $G \wr \mathbb{Z}_2$ are constructed in the following ways:

1. Each unordered pair of two non-isomorphic irreps $\sigma, \rho \in \widehat{G}$ gives rise to an irrep of $G \wr \mathbb{Z}_2$, denoted $\{\rho, \sigma\}$, with character given by:

$$\chi_{\{\rho, \sigma\}}((x, y), b) = \begin{cases} \chi_\rho(x)\chi_\sigma(y) + \chi_\rho(y)\chi_\sigma(x) & \text{if } b = 0 \\ 0 & \text{if } b = 1. \end{cases}$$

The dimension of representation $\{\rho, \sigma\}$ is equal to $\chi_{\{\rho, \sigma\}}((1, 1), 0) = 2d_\rho d_\sigma$.

2. Each irrep $\rho \in \widehat{G}$ gives rise to two irreps of $G \wr \mathbb{Z}_2$, denoted $\{\rho\}$ and $\{\rho\}'$, with characters given by:

$$\chi_{\{\rho\}}((x, y), b) = \begin{cases} \chi_\rho(x)\chi_\rho(y) & \text{if } b = 0 \\ \chi_\rho(xy) & \text{if } b = 1 \end{cases}$$

$$\chi_{\{\rho\}'}((x, y), b) = \begin{cases} \chi_\rho(x)\chi_\rho(y) & \text{if } b = 0 \\ -\chi_\rho(xy) & \text{if } b = 1. \end{cases}$$

Both representations $\{\rho\}$ and $\{\rho\}'$ have the same dimension equal d_ρ^2 .

Clearly, the number of irreps of $G \wr \mathbb{Z}$ is equal to $|\widehat{G}|^2/2 + 3|\widehat{G}|/2$, which is no more than $|\widehat{G}|^2$ as long as G has at least three irreps. Now it is easy to determine the maximal normalized characters on subgroup K .

Proposition 16. *For non-isomorphic irreps $\rho, \sigma \in \widehat{G}$,*

$$\overline{\chi}_{\{\rho, \sigma\}}(K) \leq \overline{\chi}_\rho(H_0)\overline{\chi}_\sigma(H_0).$$

For irrep $\rho \in \widehat{G}$,

$$\overline{\chi}_{\{\rho\}}(K) = \overline{\chi}_{\{\rho\}'}(K) = \max \left\{ \overline{\chi}_\rho(H_0)^2, 1/d_\rho \right\}.$$

So to bound the maximal normalized characters over K , we can turn to bounding the normalized characters on the subgroup H_0 and the dimension of an irrep of G .

6.2 Normalized characters for $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$

Recall that for the case of attacking McEliece cryptosystem, we have $G = \text{GL}_k(\mathbb{F}_q) \times S_n$ and every element $(A, P) \in H_0$ has $P \in \text{Aut}(M)$.

For $\tau \in \widehat{\text{GL}_k(\mathbb{F}_q)}$ and $\lambda \in \widehat{S}_n$, let $\tau \times \lambda$ denote the tensor product as a representation of $\text{GL}_k(\mathbb{F}_q) \times S_n$. Those tensor product representations $\tau \times \lambda$ are all irreps of $\text{GL}_k(\mathbb{F}_q) \times S_n$. Since $\overline{\chi}_{\tau \times \lambda}(S_\pi, \pi) = \overline{\chi}_\tau(S_\pi)\overline{\chi}_\lambda(\pi)$ and $\overline{\chi}_\tau(S_\pi) \leq 1$ for all $\pi \in S_n$, we have

$$\overline{\chi}_{\tau \times \lambda}(H_0) \leq \overline{\chi}_\lambda(\text{Aut}(M)).$$

As in the treatment for the symmetric group, we can bound the maximal normalized character $\overline{\chi}_\lambda(\text{Aut}(M))$ based on the minimum support of non-identity elements in $\text{Aut}(M)$, for any $\lambda \in \widehat{S}_n \setminus \Lambda_c$.

To complete bounding the maximal normalized characters on the subgroup K , it remains to bound the dimension of a representation $\tau \times \lambda$ of the group $\text{GL}_k(\mathbb{F}_q) \times S_n$ with $\lambda \in \widehat{S}_n \setminus \Lambda_c$. Since the dimension of $\tau \times \lambda$ is

$$d_{\tau \times \lambda} = d_\tau d_\lambda \geq d_\lambda,$$

we prove the following lower bound for d_λ .

Lemma 17. *Let $0 < c \leq 1/6$ be a constant. Then there is a constant $\beta > 0$ depending only on c such that for sufficiently large n and for $\lambda \in \widehat{S}_n \setminus \Lambda_c$,*

$$d_\lambda \geq e^{\beta n}.$$

Proof of Lemma 17. Consider an integer partition of n , $\lambda = (\lambda_1, \dots, \lambda_t)$, with both λ_1 and t less than $(1-c)n$. Let $\lambda' = (\lambda'_1, \dots, \lambda'_{\lambda_1})$ be the conjugate of λ , where $t = \lambda'_1 \geq \lambda'_2 \geq \dots \geq \lambda'_{\lambda_1}$ and $\sum_i \lambda'_i = n$. WLOG, assume $\lambda'_1 \leq \lambda_1$. We label all the cells of the Young diagram of shape λ as c_1, \dots, c_n , in which c_i is the i^{th} cell from the left of the first row, for $1 \leq i \leq \lambda_1$.

The dimension of λ is determined by the *hook length formula*:

$$d_\lambda = \frac{n!}{\text{Hook}(\lambda)}, \quad \text{Hook}(\lambda) = \prod_{i=1}^n \text{hook}(c_i),$$

where $\text{hook}(c_i)$ is the number of cells appearing in either the same column or the same row as the cell c_i , excluding those that are above or the the left of c_i . In particular,

$$\text{hook}(c_i) = \lambda_1 - i + \lambda'_i \quad \text{for } 1 \leq i \leq \lambda_1.$$

If $\lambda_1 \leq cn$, we have $\text{hook}(c_i) \leq t + \lambda_1 \leq 2cn$ for all i , thus

$$d_\lambda \geq \frac{n!}{(2cn)^n} \geq \frac{n^n}{e^n (2cn)^n} \geq \left(\frac{3}{e}\right)^n \geq e^{\beta n}.$$

Now we consider the case $cn < \lambda_1 < (1-c)n$. Let $\tilde{\lambda} = (\lambda_2, \dots, \lambda_t)$, this is an integer partition of $n - \lambda_1$ whose Young diagram is obtained by removing the first row of λ . Applying the hook length formula for $\tilde{\lambda}$ and the fact that $d_{\tilde{\lambda}} \geq 1$ gives us:

$$\text{Hook}(\tilde{\lambda}) = \frac{(n - \lambda_1)!}{d_{\tilde{\lambda}}} \leq (n - \lambda_1)!.$$

Then we have

$$\text{Hook}(\lambda) = \text{Hook}(\tilde{\lambda}) \prod_{i=1}^{\lambda_1} \text{hook}(c_i) \leq (n - \lambda_1)! \prod_{i=1}^{\lambda_1} \text{hook}(c_i).$$

On the other hand, we have

$$\begin{aligned} \prod_{i=1}^{\lambda_1} \text{hook}(c_i) &= \prod_{i=1}^{\lambda_1} (\lambda_1 - i + \lambda'_i) \\ &= \lambda_1! \prod_{i=1}^{\lambda_1} \left(1 + \frac{\lambda'_i - 1}{\lambda_1 - i + 1}\right) \\ &\leq \lambda_1! \exp\left(\sum_{i=1}^{\lambda_1} \frac{\lambda'_i - 1}{\lambda_1 - i + 1}\right) \quad (\text{since } 1 + x \leq e^x \text{ for all } x). \end{aligned}$$

To upper bound the exponent in the last equation, we use Chebyshev's sum inequality, which states that for any increasing sequence $a_1 \geq a_2 \geq \dots \geq a_k$ and any decreasing sequence $b_1 \leq b_2 \leq \dots \leq b_k$ or real numbers, we have $k \sum_{i=1}^k a_i b_i \leq (\sum_{i=1}^k a_i) (\sum_{i=1}^k b_i)$. Since the sequence $\{\lambda'_i - 1\}$ is increasing and the sequence $\{1/(\lambda_1 - i + 1)\}$ is decreasing, we get

$$\begin{aligned} \sum_{i=1}^{\lambda_1} \frac{\lambda'_i - 1}{\lambda_1 - i + 1} &\leq \frac{\sum_{i=1}^{\lambda_1} (\lambda'_i - 1)}{\lambda_1} \left(\sum_{i=1}^{\lambda_1} \frac{1}{\lambda_1 - i + 1}\right) \\ &= \frac{n - \lambda_1}{\lambda_1} \left(\sum_{i=1}^{\lambda_1} \frac{1}{i}\right) \leq \frac{1}{c} \left(\sum_{i=1}^{\lambda_1} \frac{1}{i}\right) \quad (\text{since } \lambda_1 > cn). \end{aligned}$$

Let r be a constant such that $1 < r/c < cn$. Bounding $1/i \leq 1$ for all $i \leq r/c$ and bounding $1/i \leq c/r$ for all $i > r/c$ yields

$$\sum_{i=1}^{\lambda_1} \frac{1}{i} \leq \frac{r}{c} + \frac{c\lambda_1}{r}.$$

Putting the pieces together, we get

$$\begin{aligned} d_\lambda &\geq \frac{n!}{(n-\lambda_1)!\lambda_1!e^{\lambda_1/r+r/c^2}} = \binom{n}{\lambda_1} e^{-\lambda_1/r-r/c^2} \\ &\geq \left(\frac{n}{\lambda_1}\right)^{\lambda_1} e^{-\lambda_1/r-r/c^2} \\ &\geq \left(\frac{e^{-1/r}}{1-c}\right)^{\lambda_1} e^{-r/c^2} \quad (\text{since } \lambda_1 < (1-c)n). \end{aligned}$$

Let $0 < \delta < \ln \frac{1}{1-c}$ be a constant and choose r large enough so that $e^{-1/r} \geq (1-c)e^\delta$. Then

$$d_\lambda \geq e^{\delta\lambda_1-r/c^2} \geq e^{\delta cn-r/c^2} \geq e^{\beta n}.$$

□

Remark The lower bound in Lemma 17 is essentially tight. To see this, consider the hook of width $(1-c)n$ and of depth cn . This hook has dimension roughly equal $\binom{n}{cn}$, which is no more than $(e/c)^{cn}$.

6.3 Proof of Theorem 8

We are ready to prove Theorem 8.

Proof of Theorem 8. To apply Theorem 4, let $0 < c < \min\{1/6, 1/4-a\}$ be a constant and S be the set of irreps of $(\mathrm{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$ of the forms $\{\tau \times \lambda, \eta \times \mu\}$, $\{\tau \times \lambda\}$, $\{\tau \times \lambda\}'$ with $\tau, \eta \in \widehat{\mathrm{GL}_k(\mathbb{F}_q)}$ and $\lambda, \mu \in \Lambda_c$, where Λ_c is mentioned in Section 4.1. Firstly, we need upper bounds for $\overline{\chi}_S(K)$, $|S|$, and d_S .

Since $\mathrm{Aut}(M)$ has minimal degree m , by Inequality (4) in Section 4.1, we have for all $\lambda \in \widehat{S}_n \setminus \Lambda_c$,

$$\overline{\chi}_\lambda(\mathrm{Aut}(M)) \leq e^{-\alpha m}.$$

Combining with Lemma 17 yields

$$\overline{\chi}_S(K) \leq \max\{e^{-2\alpha m}, e^{-\beta n}\} \leq e^{-\delta m},$$

for some constant $\delta > 0$ (we can set $\delta = \min\{2\alpha, \beta\}$).

Since $|\widehat{\mathrm{GL}_k(\mathbb{F}_q)}| \leq |\mathrm{GL}_k(\mathbb{F}_q)| \leq q^{k^2}$ and by Lemma 6, we have

$$|S| \leq |\widehat{\mathrm{GL}_k(\mathbb{F}_q)}|^2 |\Lambda_c|^2 \leq q^{2k^2} e^{O(\sqrt{n})}.$$

To bound d_S , we start with bounding the dimension of each representation in S . A representation $\{\tau \times \lambda, \eta \times \mu\}$ in S has dimension

$$2d_{\tau \times \lambda} d_{\eta \times \mu} = 2d_\tau d_\lambda d_\eta d_\mu \leq 2d_\tau d_\eta n^{2cn} \leq 2q^{k^2} n^{2cn},$$

where the first inequality follows Lemma 6. The last inequality holds because $d_\tau^2 \leq \sum_{p \in \widehat{\text{GL}_k(\mathbb{F}_q)}} d_p^2 = |\text{GL}_k(\mathbb{F}_q)|$ for any $\tau \in \widehat{\text{GL}_k(\mathbb{F}_q)}$. Similarly, a representation $\{\tau \times \lambda\}$ or $\{\tau \times \lambda\}'$ in S has dimension $d_{\tau \times \lambda}^2 \leq q^{k^2} n^{2cn}$. Hence, the maximal dimension of a representation in the set S is

$$d_S \leq 2q^{k^2} n^{2cn}.$$

Let $4a + 4c < d < 1$ be a constant and let γ_1 be any constant such that $0 < \gamma_1 < d - 4c - 4a$. Now we set the dimension threshold $D = n^{dn}$. From the upper bounds on $|S|$ and d_S , we get

$$\begin{aligned} |S| \frac{d_S^2}{D} &\leq 4q^{4k^2} e^{O(\sqrt{n})} n^{(4c-d)n} \\ &\leq 4e^{O(\sqrt{n})} n^{(4a+4c-d)n} && (\text{since } q^{k^2} \leq n^{an}) \\ &\leq n^{-\gamma_1 n} && \text{for sufficiently large } n. \end{aligned}$$

Let L be the set of all irreps of $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$ of dimension at least D . Bounding $|L|$ by the number of irreps of $(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2$, which is no more than square of the number of irreps of $\text{GL}_k(\mathbb{F}_q) \times S_n$, we have

$$|L| \leq \left| \widehat{\text{GL}_k(\mathbb{F}_q)} \right|^2 \left| \widehat{S_n} \right|^2 \leq |\text{GL}_k(\mathbb{F}_q)|^2 p(n)^2.$$

Hence, for sufficiently large n ,

$$\begin{aligned} \frac{|L|D^2}{|(\text{GL}_k(\mathbb{F}_q) \times S_n) \wr \mathbb{Z}_2|} &\leq \frac{|\text{GL}_k(\mathbb{F}_q)|^2 p(n)^2 n^{2dn}}{2 |\text{GL}_k(\mathbb{F}_q)|^2 |S_n|^2} = \frac{p(n)^2 n^{2dn}}{2(n!)^2} \\ &\leq \frac{e^{O(\sqrt{n})} n^{2dn}}{2n^{2n} e^{-2n}} \\ &\leq e^{O(n)} n^{2(d-1)n} \leq n^{-\gamma_2 n} \quad \text{so long as } \gamma_2 < 2(1-d). \end{aligned}$$

By Theorem 4, we have

$$\mathcal{D}_K \leq 4|K|^2 (e^{-\delta m} + n^{-\gamma_1 n} + n^{-\gamma_2 n}) \leq 4|K|^2 (e^{-\delta m} + n^{-\gamma n}),$$

for some constant $\gamma > 0$. This completes the proof. \square

Acknowledgments. This work was supported by the NSF under grants CCF-0829931, 0835735, and 0829917 and by the DTO under contract W911NF-04-R-0009. We are grateful to Jon Yard for helpful discussions.

References

- [1] Daniel J. Bernstein. List decoding for binary Goppa codes, 2008. Preprint.
- [2] Daniel J. Bernstein, Tanja Lange, and Christiane Peters. Attacking and defending the McEliece cryptosystem. In *PQCrypto '08: Proceedings of the 2nd International Workshop on Post-Quantum Cryptography*, pages 31–46, Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 978-3-540-88402-6.

- [3] D. Engelbert, R. Overbeck, and A. Schmidt. A summary of McEliece-type cryptosystems and their security. *J. Math. Crypt.*, 1:151–199, 2007.
- [4] William Fulton and Joe Harris. *Representation Theory - A First Course*. Springer-Verlag, New York Inc., 1991.
- [5] Michelangelo Grigni, J. Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Combinatorica*, 24(1):137–154, 2004.
- [6] Sean Hallgren, Cristopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen. Limitations of quantum coset states for graph isomorphism. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 604–617, 2006.
- [7] Julia Kempe and Aner Shalev. The hidden subgroup problem and permutation group theory. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1118–1125, 2005.
- [8] Julia Kempe, Laszlo Pyber, and Aner Shalev. Permutation groups, minimal degrees and quantum computing. *Groups, Geometry, and Dynamics*, 1(4):553–584, 2007. URL <http://xxx.lanl.gov/abs/quant-ph/0607204>.
- [9] Pierre Loidreau and Nicolas Sendrier. Weak keys in the McEliece public-key cryptosystem. *IEEE Transactions on Information Theory*, 47(3):1207–1212, 2001.
- [10] Chris Lomont. The hidden subgroup problem - review and open problems, 2004. URL [arXiv.org:quant-ph/0411037](http://arXiv.org/quant-ph/0411037).
- [11] R.J. McEliece. A public-key cryptosystem based on algebraic coding theory. *JPL DSN Progress Report*, pages 114–116, 1978.
- [12] A.J. Menezes, P.C. van Oorschot, and S.A. Vanstone. *Handbook of applied cryptography*. CRC Press, 1996.
- [13] Cristopher Moore, Alexander Russell, and Leonard J. Schulman. The symmetric group defies strong quantum Fourier sampling. *SIAM Journal of Computing*, 37:1842–1864, 2008.
- [14] E. Petrank and R.M. Roth. Is code equivalence easy to decide? *IEEE Transactions on Information Theory*, 43(5):1602 – 1604, 1997. doi: 10.1109/18.623157.
- [15] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 84–93, 2005.
- [16] Yuval Roichman. Upper bound on the characters of the symmetric groups. *Invent. Math.*, 125(3): 451–485, 1996.
- [17] John A. Ryan. Excluding some weak keys in the McEliece cryptosystem. In *Proceedings of the 8th IEEE Africon*, pages 1–5, 2007.
- [18] Peter. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26:1484–1509, 1997.

- [19] V. M. Sidelnikov and S. O. Shestakov. On insecurity of cryptosystems based on generalized Reed-Solomon codes. *Discrete Mathematics and Applications*, 2(4):439–444, 1992.
- [20] Daniel R. Simon. On the power of quantum computation. *SIAM J. Comput.*, 26(5):1474–1483, 1997.
- [21] Henning Stichtenoth. On automorphisms of geometric Goppa codes. *Journal of Algebra*, 130:113–121, 1990.
- [22] Henning Stichtenoth. *Algebraic Function Fields and Codes*. Springer, 2nd edition, 2008.
- [23] J.H van Lint. *Introduction to coding theory*. Springer-Verlag, 2nd edition, 1992.

Appendices

A Supplemental proofs for the main theorem

Proposition 18. *Let $H < G$ and g be chosen from G uniformly at random. Then for $\rho \in \widehat{G}$ and $\mathbf{b} \in B_\rho$,*

$$\mathbb{E}_g[P_{H^g}(\mathbf{b} \mid \rho)] = 1/d_\rho.$$

Proof. Schur's lemma asserts that if ρ is irreducible, the only matrices which commute with $\rho(g)$ for all g are the scalars. Hence,

$$\mathbb{E}_g[\Pi_{H^g}^\rho] = \frac{1}{|G|} \sum_{g \in G} \rho^\dagger(g) \Pi_H^\rho \rho(g) = \frac{\text{Tr}(\Pi_H^\rho)}{d_\rho} \mathbf{1}_{d_\rho},$$

which implies that

$$\mathbb{E}_g[\|\Pi_{H^g}^\rho \mathbf{b}\|^2] = \mathbb{E}_g[\langle \mathbf{b}, \Pi_{H^g}^\rho \mathbf{b} \rangle] = \langle \mathbf{b}, \mathbb{E}_g[\Pi_{H^g}^\rho] \mathbf{b} \rangle = \frac{\text{Tr}(\Pi_H^\rho)}{d_\rho}.$$

□

A.1 Proof of Lemma 14

Proof of Lemma 14. Let L_σ be the subspace of $\rho \otimes \rho^*$ consisting of all copies of σ . Since B_ρ is orthonormal, the vectors $\{\mathbf{b} \otimes \mathbf{b}^* \mid \mathbf{b} \in B_\rho\}$ are mutually orthogonal in $\rho \otimes \rho^*$. Thus,

$$\sum_{\mathbf{b} \in B_\rho} \left\| \Pi_\sigma^{\rho \otimes \rho^*}(\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 \leq \dim L_\sigma.$$

Note that $\dim L_\sigma$ is equal to d_σ times the multiplicity of σ in $\rho \otimes \rho^*$. On the other hand, we have

$$\begin{aligned} \text{multiplicity of } \sigma \text{ in } \rho \otimes \rho^* &= \langle \chi_\sigma, \chi_\rho \chi_{\rho^*} \rangle = \langle \chi_\sigma \chi_\rho, \chi_{\rho^*} \rangle \\ &= \text{multiplicity of } \rho^* \text{ in } \sigma \otimes \rho \\ &\leq \frac{\dim(\sigma \otimes \rho)}{\dim \rho^*} = d_\sigma, \end{aligned}$$

Hence,

$$\sum_{\mathbf{b} \in B_\rho} \left\| \Pi_\sigma^{\rho \otimes \rho^*}(\mathbf{b} \otimes \mathbf{b}^*) \right\|^2 \leq d_\sigma^2.$$

□

B Rational Goppa codes

This part summarizes definitions and key properties of rational Goppa codes that would be useful in our analysis. Following Stichtenoth [21], we shall describe Goppa codes in terms of algebraic function fields instead of algebraic curves. A complete treatment for this subject can be found in [22].

A rational function field over \mathbb{F}_q is a field extension $\mathbb{F}_q(x)/\mathbb{F}_q$ for some x transcendental over \mathbb{F}_q . Each element $z \in \mathbb{F}_q(x)$ can be viewed as a function whose evaluation at a base field element $a \in \mathbb{F}_q$ is determined as follows: write $z = f(x)/g(x)$ for some polynomials $f(x), g(x) \in \mathbb{F}_q[x]$, then

$$z(a) = \begin{cases} \frac{f(a)}{g(a)} \in \mathbb{F}_q & \text{if } g(a) \neq 0 \\ \infty & \text{if } g(a) = 0. \end{cases}$$

A Riemann-Roch space⁴ in the rational function field $\mathbb{F}_q(x)/\mathbb{F}_q$ is a subset of \mathbb{F}_q of the form

$$\mathcal{L}(r, g, h) = \left\{ \frac{f(x)g(x)}{h(x)} \mid f(x) \in \mathbb{F}_q[x], \deg f(x) \leq r \right\}$$

for some nonzero polynomials $g(x), h(x) \in K[x]$ and some integer r . Note that $\mathcal{L}(r, g, h)$ is a vector space of dimension $r + 1$ over \mathbb{F}_q .

Definition. (A special case of Definition 2.2.1 in [22]) Let $g(x), h(x) \in \mathbb{F}_q[x]$ be nonzero coprime polynomials, and let $r < n$ a nonnegative integer. Let $\gamma_1, \dots, \gamma_n$ be n distinct elements in the field⁵ \mathbb{F}_q such that $g(\gamma_i) \neq 0$ and $h(\gamma_i) \neq 0$ for all i . Then a rational Goppa code associated with g, h and γ_i 's is defined by

$$\mathcal{C}(\gamma_1, \dots, \gamma_n, r, g, h) \stackrel{\text{def}}{=} \{(z(\gamma_1), \dots, z(\gamma_n)) \mid z \in \mathcal{L}(r, g, h)\} \subset \mathbb{F}_q^n.$$

Remark A classical binary Goppa code can be obtained by setting $q = 2^m$, $r = n - \deg g(x) - 1$, and $h(x) = \sum_{j=1}^n \prod_{i \neq j} (x - \gamma_i)$ and then intersecting the code $\mathcal{C}(\gamma_1, \dots, \gamma_n, r, g, h)$ with the vector space \mathbb{F}_2^n (see [1]). Generalized Reed-Solomon codes are a special case of rational Goppa codes in which the polynomials $g(x)$ and $h(x)$ are both constants.

Theorem 19. (A special case of Corollary 2.2.3 in [22]) The code defined in Definition B is an $[n, k, d]$ -linear code over \mathbb{F}_q with dimension $k = r + 1$ and minimum distance $d \geq n - r$. Consequentially, this code can correct at least $(n - r - 1)/2$ errors.

The rational Goppa code $\mathcal{C}(\gamma_1, \dots, \gamma_n, r, g, h)$ has a generator matrix:

$$M_0 = \begin{pmatrix} \frac{g(\gamma_1)}{h(\gamma_1)} & \cdots & \frac{g(\gamma_n)}{h(\gamma_n)} \\ \gamma_1 \frac{g(\gamma_1)}{h(\gamma_1)} & \cdots & \gamma_n \frac{g(\gamma_n)}{h(\gamma_n)} \\ \vdots & \ddots & \vdots \\ \gamma_1^r \frac{g(\gamma_1)}{h(\gamma_1)} & \cdots & \gamma_n^r \frac{g(\gamma_n)}{h(\gamma_n)} \end{pmatrix}.$$

Proposition 20. The matrix M_0 has full rank, that is, its column rank equals $r + 1$. Hence, every generator matrix of a rational Goppa code has full rank.

⁴In terms of algebraic function fields, a Riemann-Roch space is defined in the association with a divisor of the function field F/K , where a divisor is a finite sum $\sum_i n_i P_i$ with $n_i \in \mathbb{Z}$ and P_i 's being places of the function field. In the rational function field $K(x)/K$, we can show that every divisor can be written as $rP_\infty + (z)$ for some integer r and some nonzero $z \in K(x)$, where P_∞ is the infinite place (defined in [22, pg. 9]), and (z) is the principal divisor of z . The space $\mathcal{L}(r, g, h)$ is indeed the Riemann-Roch space associated with the divisor $rP_\infty + (z)$ with $z = h(x)/g(x)$.

⁵In the case $r = \deg h(x) - \deg g(x)$, one can choose one of the points P_i 's to be ∞ . However, we rule out this case to keep the discussion simple.

Proof. It suffices to show that the first $r+1$ columns of M_0 are linearly independent. Equivalently, we show that the matrix N_0 below has nonzero determinant:

$$N_0 = \begin{pmatrix} \frac{g(\gamma_1)}{h(\gamma_1)} & \cdots & \frac{g(\gamma_{r+1})}{h(\gamma_{r+1})} \\ \gamma_1 \frac{g(\gamma_1)}{h(\gamma_1)} & \cdots & \gamma_{r+1} \frac{g(\gamma_{r+1})}{h(\gamma_{r+1})} \\ \vdots & \ddots & \vdots \\ \gamma_1^r \frac{g(\gamma_1)}{h(\gamma_1)} & \cdots & \gamma_{r+1}^r \frac{g(\gamma_{r+1})}{h(\gamma_{r+1})} \end{pmatrix} = \begin{pmatrix} 1 & \cdots & 1 \\ \gamma_1 & \cdots & \gamma_{r+1} \\ \vdots & \ddots & \vdots \\ \gamma_1^r & \cdots & \gamma_{r+1}^r \end{pmatrix} \begin{pmatrix} \frac{g(\gamma_1)}{h(\gamma_1)} & & \\ & \ddots & \\ & & \frac{g(\gamma_{r+1})}{h(\gamma_{r+1})} \end{pmatrix}.$$

The first matrix in the above product is a Vandermonde matrix, which has nonzero determinant because γ_i 's are distinct. The second matrix also has nonzero determinant because $g(\gamma_i) \neq 0$ for all i . Hence, N_0 has nonzero determinant. \square

An important property of rational Goppa codes is that in general their automorphisms are induced by projective transformations of the projective line. We will make this precise below.

Definition. (See [23, pg. 53]) Let C be a code of length n . An *automorphism* of C is a permutation $\pi \in S_n$ which maps every word in C to a word in C by acting on the positions of the codewords. The set of all automorphisms of C forms a group called the **automorphism group** of C .

In particular, an automorphism of $\mathcal{C}(\gamma_1, \dots, \gamma_n, r, g, h)$ is a permutation $\pi \in S_n$ such that

$$\mathcal{C}(\gamma_1, \dots, \gamma_n, r, g, h) = \mathcal{C}(\gamma_{\pi(1)}, \dots, \gamma_{\pi(n)}, r, g, h).$$

Remark Suppose M is a generator matrix for an $[n, k]$ -linear code C over \mathbb{F}_q . Then a permutation $\pi \in S_n$ is an automorphism of C if and only if there is an invertible matrix $A \in \text{GL}_k(\mathbb{F}_q)$ such that $AMP_\pi = P_\pi$, where P_π denotes the permutation matrix corresponding to π . If M has full rank, there is exactly one such matrix A for each automorphism π of C .

Theorem 21 (Stichtenoth [21]). *Suppose $1 \leq r \leq n-3$. Then the automorphism group of the rational Goppa code $\mathcal{C}(\gamma_1, \dots, \gamma_n, r, g, h)$ is isomorphic to a subgroup of $\text{Aut}(\mathbb{F}_q(x)/\mathbb{F}_q)$.*

Fact. The automorphism group of the rational function field $\mathbb{F}_q(x)/\mathbb{F}_q$ is isomorphic to the projective linear group over \mathbb{F}_q . In notations, $\text{Aut}(\mathbb{F}_q(x)/\mathbb{F}_q) \simeq \text{PGL}_2(\mathbb{F}_q)$.

Let $C = \mathcal{C}(\gamma_1, \dots, \gamma_n, r, g, h)$ be a rational Goppa code. To give an intuition for how the automorphism group of C is embedded in $\text{PGL}_2(\mathbb{F}_q)$, consider a transformation $\sigma \in \text{PGL}_2(\mathbb{F}_q)$ and view each element $a \in \mathbb{F}_q$ as the projective line $[a : 1]$ (the point at infinity is written as $[1 : 0]$). Suppose σ transforms $[a : 1]$ to the projective line $[b : 1]$, then we shall write $\sigma a = b$. If σ transforms each line $[\gamma_i : 1]$ to some line $[\gamma_j : 1]$, then σ induces another rational Goppa code:

$$\mathcal{C}(\sigma\gamma_1, \dots, \sigma\gamma_n, r, g, h).$$

If, further, $\mathcal{C}(\sigma\gamma_1, \dots, \sigma\gamma_n, r, g, h)$ equals the original code C , then σ induces an automorphism of C . Stichtenoth's theorem establishes that every automorphism of C is induced by such a transformation in $\text{PGL}_2(\mathbb{F}_q)$.

C Supplemental proofs for $\text{GL}_2(\mathbb{F}_q)$

C.1 Irreducible representations of $\text{GL}_2(\mathbb{F}_q)$

In this part, we will first present a preliminary background on the structure of $\text{GL}_2(\mathbb{F}_q)$ followed by description of its irreps. We refer readers to [4, §5.2] for the missing technical details in this part.

Viewing $\text{GL}_2(\mathbb{F}_q)$ as the group of all \mathbb{F}_q -linear invertible endomorphisms of the quadratic extension \mathbb{F}_{q^2} of \mathbb{F}_q , we have a large subgroup of $\text{GL}_2(\mathbb{F}_q)$ that is isomorphic to $\mathbb{F}_{q^2}^*$ via the identification:

$$\{f_\xi \mid \xi \in \mathbb{F}_{q^2}^*\} \simeq \mathbb{F}_{q^2}^*, \quad f_\xi \leftrightarrow \xi$$

where $f_\xi : \mathbb{F}_{q^2} \rightarrow \mathbb{F}_{q^2}$ is the \mathbb{F}_q -linear map given by $f_\xi(v) = \xi v$ for all $v \in \mathbb{F}_{q^2}$.

To turn each map f_ξ into a matrix form, we fix a basis $\{1, \gamma\}$ of \mathbb{F}_{q^2} as a vector space over \mathbb{F}_q . For each $\xi \in \mathbb{F}_{q^2}$, writing $\xi = \xi_{x,y} = x + \gamma y$ for some $x, y \in \mathbb{F}_q$, then the map f_ξ corresponds to the matrix $\begin{pmatrix} x & \gamma^2 y \\ y & x \end{pmatrix}$, since $f_\xi(1) = x + \gamma y$ and $f_\xi(\gamma) = \gamma^2 y + \gamma x$. Hence, we can rewrite the above identification as

$$\left\{ \begin{pmatrix} x & \gamma^2 y \\ y & x \end{pmatrix} \mid x, y \in \mathbb{F}_q, x \neq 0 \text{ or } y \neq 0 \right\} \simeq \mathbb{F}_{q^2}^*, \quad \begin{pmatrix} x & \gamma^2 y \\ y & x \end{pmatrix} \leftrightarrow \xi_{x,y} = x + \gamma y.$$

For example, if q is odd, choose a generator ε of \mathbb{F}_q^* , then ε must be non-square in \mathbb{F}_q , which implies that $\{1, \sqrt{\varepsilon}\}$ form a basis of \mathbb{F}_{q^2} as a vector space over \mathbb{F}_q . In such a case, we can define $\xi_{x,y} = x + y\sqrt{\varepsilon}$.

Conjugacy classes. Group $\text{GL}_2(\mathbb{F}_q)$ has four types of conjugacy classes in Table C.1, with representatives described as follows:

$$a_x = \begin{pmatrix} x & 0 \\ 0 & x \end{pmatrix} \quad b_x = \begin{pmatrix} x & 1 \\ 0 & x \end{pmatrix} \quad c_{x,y} = \begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix} \quad d_{x,y} = \begin{pmatrix} x & \gamma^2 y \\ y & x \end{pmatrix}$$

class	$[a_x]$ $x \in \mathbb{F}_q^*$	$[b_x]$ $x \in \mathbb{F}_q^*$	$[c_{x,y}] = [c_{y,x}]$ $x, y \in \mathbb{F}_q^*, x \neq y$	$[d_{x,y}] = [d_{x,-y}]$ $x \in \mathbb{F}_q, y \in \mathbb{F}_q^*$
class size	1	$q^2 - 1$	$q^2 + q$	$q^2 - q$
no. of classes	$q - 1$	$q - 1$	$\frac{(q-1)(q-2)}{2}$	$\frac{q(q-1)}{2}$

Table 1: Conjugacy classes of $\text{GL}_2(\mathbb{F}_q)$, where $[g]$ denotes the class of representative g .

There are $q^2 - 1$ conjugacy classes, hence there are exactly $q^2 - 1$ irreps of $\text{GL}_2(\mathbb{F}_q)$. We shall briefly describe below how to construct all those representations.

Linear representations. For each character $\alpha : \mathbb{F}_q^* \rightarrow \mathbb{C}^*$ of the cyclic group \mathbb{F}_q^* , we have a one-dimensional representation U_α of $\text{GL}_2(\mathbb{F}_q)$ defined by:

$$U_\alpha(g) = \alpha(\det(g)) \quad \forall g \in \text{GL}(2, q).$$

To compute $U_\alpha(d_{x,y})$, we shall use the following fact:

$$\det \begin{pmatrix} x & \gamma^2 y \\ y & x \end{pmatrix} = \text{Norm}_{\mathbb{F}_{q^2}/\mathbb{F}_q}(\xi_{x,y}) = \xi_{x,y} \cdot \xi_{x,y}^q = \xi_{x,y}^{q+1}.$$

Recall that there are $q-1$ characters of $\mathbb{F}_q^* = \langle \varepsilon \rangle$ corresponding to $q-1$ places where the generator ε can be sent to. The linear representation U_{α_0} , where α_0 is the character sending ε to 1, is indeed the trivial representation, denoted U .

Irreducible representations by action on $\mathbb{P}^1(\mathbb{F}_q)$. $\text{GL}_2(\mathbb{F}_q)$ acts transitively on the projective line $\mathbb{P}^1(\mathbb{F}_q)$ in the natural way:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot [x : y] = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = [ax + by : cx + dy],$$

in which the stabilizer of the infinite point $[1 : 0]$ is the Borel subgroup B :

$$B = \left\{ \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mid a, d \in \mathbb{F}_q^*, b \in \mathbb{F}_q \right\}.$$

The permutation representation of $\text{GL}_2(\mathbb{F}_q)$ given by this action on $\mathbb{P}^1(\mathbb{F}_q)$ has dimension $q+1$ and decomposes into the trivial representation U and a q -dimensional representation V . The character of V is given as follows:

$$\chi_V(a_x) = q \quad \chi_V(b_x) = 0 \quad \chi_V(c_{x,y}) = 1 \quad \chi_V(d_{x,y}) = -1.$$

By checking $\langle \chi_V, \chi_V \rangle = 1$, we see that V is irreducible. Hence, for each of the $q-1$ characters α of \mathbb{F}_q^* , we have a q -dimensional irrep $V_\alpha = V \otimes U_\alpha$. Note that $V = V \otimes U$.

Irreducible representations induced from Borel subgroup B . For each pair of characters α, β of \mathbb{F}_q^* , there is a character of the subgroup B :

$$\phi_{\alpha,\beta} : B \rightarrow \mathbb{C}^* \quad \text{by} \quad \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \mapsto \alpha(a)\beta(d).$$

In other words, $\phi_{\alpha,\beta}$ is a one-dimensional representation of subgroup B . Let $W_{\alpha,\beta}$ be the representation of $\text{GL}_2(\mathbb{F}_q)$ induced by $\phi_{\alpha,\beta}$. By computing characters, we have

- $W_{\alpha,\beta} = W_{\beta,\alpha}$,
- $W_{\alpha,\alpha} = U_\alpha \oplus V_\alpha$, and
- $W_{\alpha,\beta}$ is irreducible for $\alpha \neq \beta$. Each of these representations has dimension equal the index of B in $\text{GL}_2(\mathbb{F}_q)$, i.e., $[\text{GL}(2, q) : B] = q+1$.

There are $((q-1)^2 - (q-1))/2 = (q-1)(q-2)/2$ distinct irreps of this type.

Irreducible representations by characters of $\mathbb{F}_{q^2}^*$. Let $\varphi : \mathbb{F}_{q^2}^* \rightarrow \mathbb{C}^*$ be a character of the cyclic group $\mathbb{F}_{q^2}^*$. Since $\mathbb{F}_{q^2}^*$ can be viewed as a subgroup of $\text{GL}_2(\mathbb{F}_q)$, we have the induced representation $\text{Ind} \varphi$, which is not irreducible. However, it gives us a $(q-1)$ -dimensional irrep with character given by

$$\chi_\varphi = \chi_{V \otimes W_{\alpha,1}} - \chi_{W_{\alpha,1}} - \chi_{\text{Ind} \varphi} \quad \text{if } \varphi|_{\mathbb{F}_q^*} = \alpha.$$

Note that $\text{Ind} \varphi \simeq \text{Ind} \varphi^q$, thus $X_\varphi \simeq X_{\varphi^q}$. So, the characters φ of $\mathbb{F}_{q^2}^*$ with $\varphi \neq \varphi^q$ give a rise to the $\frac{1}{2}q(q-1)$ remaining irreps of $\text{GL}_2(\mathbb{F}_q)$.

A summary of all irreducible characters of $\text{GL}_2(\mathbb{F}_q)$ is given in Table C.1 below.

ρ	d_ρ	$\chi_\rho(a_x)$	$\chi_\rho(b_x)$	$\chi_\rho(c_{x,y})$	$\chi_\rho(d_{x,y})$
U_α	1	$\alpha(x^2)$	$\alpha(x^2)$	$\alpha(xy)$	$\alpha(\xi_{x,y}^{q+1})$
V_α	q	$q\alpha(x^2)$	0	$\alpha(xy)$	$-\alpha(\xi_{x,y}^{q+1})$
$W_{\alpha,\beta} (\alpha \neq \beta)$	$q+1$	$(q+1)\alpha(x)\beta(x)$	$\alpha(x)\beta(x)$	$\alpha(x)\beta(y) + \alpha(y)\beta(x)$	0
X_φ	$q-1$	$(q-1)\varphi(x)$	$-\varphi(x)$	0	$-(\varphi(\xi_{x,y}) + \varphi(\xi_{x,y}^q))$

Table 2: Character table of $\text{GL}_2(\mathbb{F}_q)$, where α, β are characters of \mathbb{F}_q^* , and φ is a character of $\mathbb{F}_{q^2}^*$ with $\varphi^q \neq \varphi$, and $d_\rho = \chi_\rho(a_1)$ is the dimension of ρ .

C.2 Proof of Lemma 11

In the remaining of this section, we devote to prove Lemma 11, which states that there are at most two linear representations appearing in the decomposition of $\rho \otimes \rho^*$, for any irrep ρ of $\text{GL}_2(\mathbb{F}_q)$. Obviously, if ρ is linear then $\rho \otimes \rho^*$ is the trivial representation. Therefore, we shall only consider the cases where ρ is non-linear.

Recall that the multiplicity of U_α in $\rho \otimes \rho^*$ is given by

$$\langle \chi_{\rho \otimes \rho^*}, \chi_{U_\alpha} \rangle = \frac{1}{|G|} \sum_{g \in G} |\chi_\rho(g)|^2 \chi_{U_\alpha}(g) = \frac{1}{|G|} (A(\rho, \alpha) + B(\rho, \alpha) + C(\rho, \alpha) + D(\rho, \alpha)),$$

where $A(\rho, \alpha), B(\rho, \alpha), C(\rho, \alpha), D(\rho, \alpha)$ are the sum of $|\chi_\rho(g)|^2 \chi_{U_\alpha}(g)$ over all element g in the conjugacy classes with representatives of the form $a_x, b_x, c_{x,y}$ and $d_{x,y}$, respectively. That is, from the description of conjugacy classes in Table C.1,

$$\begin{aligned} A(\rho, \alpha) &= \sum_{x \in \mathbb{F}_q^*} |\chi_\rho(a_x)|^2 \chi_{U_\alpha}(a_x) \\ B(\rho, \alpha) &= (q^2 - 1) \sum_{x \in \mathbb{F}_q^*} |\chi_\rho(b_x)|^2 \chi_{U_\alpha}(b_x) \\ C(\rho, \alpha) &= \frac{1}{2}(q^2 + q) \sum_{x,y \in \mathbb{F}_q^*, x \neq y} |\chi_\rho(c_{x,y})|^2 \chi_{U_\alpha}(c_{x,y}) \\ D(\rho, \alpha) &= \frac{1}{2}(q^2 - q) \sum_{x,y \in \mathbb{F}_q, y \neq 0} |\chi_\rho(d_{x,y})|^2 \chi_{U_\alpha}(d_{x,y}). \end{aligned}$$

Our goal below will be to show that $\langle \chi_{\rho \otimes \rho^*}, \chi_{U_\alpha} \rangle = 0$ for all but two linear representations U_α and for all non-linear irrep ρ of $\text{GL}_2(\mathbb{F}_q)$. We begin with the following lemma.

Lemma 22. *Let F be a finite field and $\phi : F^\times \rightarrow \mathbb{C}^*$ be a non-trivial character of the cyclic group F^\times , i.e., $\phi(x) \neq 1$ for some x . Then $\sum_{x \in F^\times} \phi(x) = 0$.*

Proof. Let n be the order of F^\times and let τ be a generator of F^\times . Then $\tau^n = 1$ which implies $\phi(\tau)^n = 1$. Since ϕ is non-trivial, we must have $\phi(\tau) \neq 1$. Hence,

$$\sum_{x \in F^\times} \phi(x) = \sum_{k=0}^{n-1} \phi(\tau^k) = \sum_{k=0}^{n-1} \phi(\tau)^k = \frac{\phi(\tau)^n - 1}{\phi(\tau) - 1} = 0.$$

□

Note that for any character α of \mathbb{F}_q^* , the map $\alpha^2 : \mathbb{F}_q^* \rightarrow \mathbb{C}^*$ defined by $\alpha^2(x) = \alpha(x^2)$ is also a character of \mathbb{F}_q^* . Hence, we have the following direct corollaries of Lemma 22.

Corollary 23. *Let α be a character of \mathbb{F}_q^* such that α^2 is non-trivial. Then $\sum_{x \in \mathbb{F}_q^*} \alpha(x^2) = 0$.*

Corollary 24. *Let ρ be an irrep of $\mathrm{GL}_2(\mathbb{F}_q)$ and let α be a character of \mathbb{F}_q^* such that α^2 is non-trivial. Then we always have $A(\rho, \alpha) = B(\rho, \alpha) = 0$.*

Proof. Observe that $|\chi_\rho(a_x)|$ and $|\chi_\rho(b_x)|$ do not depend on x , and $\chi_{U_\alpha}(a_x) = \chi_{U_\alpha}(b_x) = \alpha(x^2)$. Hence, to show $A(\rho, \alpha) = B(\rho, \alpha) = 0$, it suffices to use the fact that $\sum_{x \in \mathbb{F}_q^*} \alpha(x^2) = 0$. \square

Remark There are at most two characters α of \mathbb{F}_q^* such that α^2 is trivial. They are the trivial one, and the one that maps $\varepsilon \rightarrow \omega^{\frac{q-1}{2}}$ if q is odd, where $\omega = e^{\frac{2\pi i}{q-1}}$ is a primitive $(q-1)^{\text{th}}$ root of unity, and ε is a chosen generator of the cyclic group \mathbb{F}_q^* . To see this, suppose $\alpha(\varepsilon) = \omega^k$, for some $k \in \{0, 1, \dots, q-2\}$. If $\alpha(\varepsilon)^2 = 1$, then $\omega^{2k} = 1$, which implies $q-1 \mid 2k$ because ω has order $q-1$. Hence $2k \in \{0, q-1\}$.

With this remark, Lemma 11 will immediately follow Lemma 25 below.

Lemma 25. *Let ρ be a non-linear irrep of $\mathrm{GL}_2(\mathbb{F}_q)$ and let α be a character of \mathbb{F}_q^* such that α^2 is trivial. Then U_α does not appear in the decomposition of $\rho \otimes \rho^*$.*

Proof. We will prove case by case of ρ that $C(\rho, \alpha) = D(\rho, \alpha) = 0$, which, together with Corollary 24, will complete the proof for the lemma.

Case $\rho = W_{\beta, \beta'}$. For this case, as $|\chi_{W_{\beta, \beta'}}(d_{x,y})| = 0$, we only need to show $C(W_{\beta, \beta'}, \alpha) = 0$. Considering $x, y \in \mathbb{F}_q^*$ with $x \neq y$ and letting $z = x^{-1}y \neq 1$, we have

$$\begin{aligned} |\chi_{W_{\beta, \beta'}}(c_{x,y})|^2 &= [\beta(x)\beta'(y) + \beta(y)\beta'(x)][\beta(x^{-1})\beta'(y^{-1}) + \beta(y^{-1})\beta'(x^{-1})] \\ &= 2 + \beta(xy^{-1})\beta'(yx^{-1}) + \beta(yx^{-1})\beta'(xy^{-1}) \\ &= 2 + \beta(z^{-1})\beta'(z) + \beta(z)\beta'(z^{-1}) \end{aligned}$$

This means $|\chi_{W_{\beta, \beta'}}(c_{x,y})|^2$ only depends on $z = x^{-1}y$. Now let $\gamma(z) = |\chi_{W_{\beta, \beta'}}(c_{x,y})|^2 \alpha(z)$, we have

$$|\chi_{W_{\beta, \beta'}}(c_{x,y})|^2 \chi_{U_\alpha}(c_{x,y}) = |\chi_{W_{\beta, \beta'}}(c_{x,y})|^2 \alpha(x^2 z) = \gamma(z) \alpha(x^2).$$

Hence,

$$\begin{aligned} \sum_{x,y \in \mathbb{F}_q^*, x \neq y} |\chi_\rho(c_{x,y})|^2 \chi_{U_\alpha}(c_{x,y}) &= \sum_{x,z \in \mathbb{F}_q^*, z \neq 1} \gamma(z) \alpha(x^2) \\ &= \left(\sum_{x \in \mathbb{F}_q^*} \alpha(x^2) \right) \left(\sum_{z \in \mathbb{F}_q^*, z \neq 1} \gamma(z) \right) = 0 \end{aligned}$$

by Corollary 23, completing the proof for the case $\rho = W_{\beta, \beta'}$.

Case $\rho = V_\beta$. Since $|\chi_{V_\beta}(c_{x,y})| = 1$ and $\chi_{U_\alpha}(c_{x,y}) = \alpha(xy) = \alpha(x)\alpha(y)$,

$$\sum_{x,y \in \mathbb{F}_q^*, x \neq y} |\chi_{V_\beta}(c_{x,y})|^2 \chi_{U_\alpha}(c_{x,y}) = \sum_{x,y \in \mathbb{F}_q^*, x \neq y} \alpha(x)\alpha(y) = \left(\sum_{x \in \mathbb{F}_q^*} \alpha(x) \right)^2 - \sum_{x \in \mathbb{F}_q^*} \alpha(x^2) = 0$$

by Lemma 22 and Corollary 23. This shows $C(V_\beta, \alpha) = 0$.

Now we are going to show that $D(V_\beta, \alpha) = 0$, or equivalently, $\sum_{x,y \in \mathbb{F}_q, y \neq 0} \alpha(\xi_{x,y}^{q+1}) = 0$. We have

$$\sum_{\xi \in \mathbb{F}_{q^2}^*} \alpha(\xi^{q+1}) = \sum_{x,y \in \mathbb{F}_q, y \neq 0} \alpha(\xi_{x,y}^{q+1}) + \sum_{x \in \mathbb{F}_q^*} \alpha(\xi_{x,0}^{q+1}) = \sum_{x,y \in \mathbb{F}_q, y \neq 0} \alpha(\xi_{x,y}^{q+1}).$$

where in the last equality, we apply Corollary 23 and the fact that $\xi_{x,0}^{q+1} = x^{q+1} = x^2$ for all $x \in \mathbb{F}_q^*$.

Consider the map $\phi : \mathbb{F}_{q^2}^* \rightarrow \mathbb{C}^*$ given by $\phi(\xi) = \alpha(\xi^{q+1})$. Clearly, ϕ is a character of $\mathbb{F}_{q^2}^*$. Since α^2 is non-trivial and $\alpha^2(x) = \alpha(x^2) = \alpha(x^{q+1}) = \phi(x)$ for all $x \in \mathbb{F}_q^*$, the map ϕ is also non-trivial. By Lemma 22, we have $\sum_{\xi \in \mathbb{F}_{q^2}^*} \alpha(\xi^{q+1}) = 0$, which implies $D(V_\beta, \alpha) = 0$.

Case $\rho = X_\varphi$. As it is clear from the character table of $\text{GL}_2(\mathbb{F}_q)$ that $C(X_\varphi, \alpha) = 0$, it remains to show $D(X_\varphi, \alpha) = 0$, or equivalently, $D_0 \stackrel{\text{def}}{=} \sum_{x,y \in \mathbb{F}_q, y \neq 0} |\varphi(\xi_{x,y}) + \varphi(\xi_{x,y}^q)|^2 \alpha(\xi_{x,y}^{q+1}) = 0$. We have

$$D_0 = \underbrace{\sum_{\xi \in \mathbb{F}_{q^2}^*} |\varphi(\xi) + \varphi(\xi^q)|^2 \alpha(\xi^{q+1})}_{D_1} - \underbrace{\sum_{x \in \mathbb{F}_q^*} |\varphi(\xi_{x,0}) + \varphi(\xi_{x,0}^q)|^2 \alpha(\xi_{x,0}^{q+1})}_{D_2}.$$

For $\xi \in \mathbb{F}_{q^2}^*$, we have

$$|\varphi(\xi) + \varphi(\xi^q)|^2 = (\varphi(\xi) + \varphi(\xi^q))(\varphi(\xi)^{-1} + \varphi(\xi^q)^{-1}) = 2 + \varphi(\xi^{q-1}) + \varphi(\xi^{1-q}).$$

Hence, since $x^{q-1} = 1$ for all $x \in \mathbb{F}_q^*$ and by Corollary 23,

$$D_2 = \sum_{x \in \mathbb{F}_q^*} (2 + \varphi(x^{q-1}) + \varphi(x^{1-q})) \alpha(x^{q+1}) = 3 \sum_{x \in \mathbb{F}_q^*} \alpha(x^2) = 0.$$

The last thing we want to show is that $D_1 = 0$. Consider the map $\phi : \mathbb{F}_{q^2}^* \rightarrow \mathbb{C}^*$ given by $\phi(\xi) = \varphi(\xi^{q-1})\alpha(\xi^{q+1})$, which is apparently a character of $\mathbb{F}_{q^2}^*$. We shall see that it is non-trivial. Let ω be a generator of $\mathbb{F}_{q^2}^*$. Since $\omega^{q^2-1} = 1$, we have $\phi(\omega^{q+1}) = \alpha(\omega^{(q+1)^2}) = \alpha(\omega^{2(q+1)}) = \alpha^2(\omega^{q+1})$. On the other hand, ω^{q+1} is a generator for \mathbb{F}_q^* , because $\omega^{k(q+1)}$ with $k = 0, 1, \dots, q-2$ are distinct, and $\omega^{(q-1)(q+1)} = 1$. Hence, if $\phi(\omega^{q+1}) = 1$, then $\alpha^2(x) = 1$ for all $x \in \mathbb{F}_q^*$. But since α^2 is non-trivial, we must have $\phi(\omega^{q+1}) \neq 1$, which means ϕ is non-trivial. Applying Lemma 22, we get $\sum_{\xi \in \mathbb{F}_{q^2}^*} \varphi(\xi^{q-1})\alpha(\xi^{q+1}) = 0$. Similarly, we also have $\sum_{\xi \in \mathbb{F}_{q^2}^*} \varphi(\xi^{1-q})\alpha(\xi^{q+1}) = 0$. Combining with the fact that $\sum_{\xi \in \mathbb{F}_{q^2}^*} \alpha(\xi^{q+1}) = 0$, which has been proved in the previous case, we have shown $D_1 = 0$, completing the proof. \square